

Reality from One Distinction: A Forcing Chain for Spacetime, Time, and the Fundamental Constants

Jonathan Washburn*

April 2026

Abstract

We exhibit a forcing chain from a single propositional commitment, the existence of a non-trivial distinction on some carrier, to the architecture of physical reality: the Lawvere natural-number object, a discrete recognition lattice, the spatial dimension $D = 3$, the golden ratio φ as the unique self-similar scaling, the Lorentzian $(1, 3)$ signature with its light cone, and the constants c , \hbar , and G as φ -powers in recognition-native units. Each link of the chain is a theorem; no link is a free parameter or empirical input. The chain is fully machine-verified in Lean 4: the master theorem `reality_from_one_distinction` compiles cleanly with zero `sorry` declarations and no theory-specific axioms, depending only on the standard kernel set `{propext, Classical.choice, Quot.sound}`. The present paper translates the chain into mathematical prose for a foundations-of-physics audience, places it in the context of prior foundational programs (Wheeler, Tegmark, Connes, Rovelli, Smolin, Bilson-Thompson, Markopoulou), presents the master theorem and its bundled *reality certificate* in math notation, and exhibits two worked numerical examples (α^{-1} inside CODATA, η_B at the Planck 2018 central value) and one mass-ladder example (the electron mass). We are explicit, throughout, about which claims are theorems, which are empirical confirmations of theorem-grade predictions, and which remain hypotheses with named falsifiers. The framework's interpretive freedom is small and we identify it precisely; most of the chain's correspondence with observation is a sequence of successful predictions, not a stack of model commitments.

Keywords: foundations of physics; functional equation; recognition; Aristotelian conditions; Lawvere natural-number object; dimension forcing; golden ratio; spacetime emergence; fundamental constants; Lean 4; machine verification.

MSC 2020: 03A05, 03B22, 18A40, 39B22, 51A05, 81P15, 83A05.

Formalization. Every theorem stated in this paper is formalized in Lean 4 + Mathlib. The complete machine-checked chain is at <https://github.com/jonwashburn/shape-of-logic>. The query `#print axioms reality_from_one_distinction` returns the standard kernel set `{propext, Classical.choice, Quot.sound}` with no theory-specific axioms. Appendix B reproduces the audit and the build receipt.

Contents

1	Introduction	2
2	The Absolute Floor	6

*Recognition Physics Institute, Austin, TX 78701, USA. jon@recognitionphysics.org.

3	Aristotelian Conditions on a Comparison Operator	7
4	Magnitude of Mismatch and Composition Consistency	8
5	The Recognition Composition Law and Branch Selection	9
6	The Cost Functional Equation: J-Uniqueness via Aczél	11
7	Arithmetic from Logic: The Lawvere NNO	12
8	Time as the Forced Orbit	14
9	Recognition Lattices from Finite-Resolution Recognizers	14
10	Dimension Forcing: $D = 3$	16
11	The Golden Ratio: Self-Similar Scaling	17
12	Spacetime Emergence: Lorentzian $(1, 3)$	18
13	Constants: c, \hbar, and G as φ-powers	20
14	The Master Theorem: Reality from One Distinction	22
15	Audit, Honest Scope, Falsification, Open Frontiers	24
A	Glossary of Definitions	27
B	Machine-Checked Verification and Reproducibility	28

1 Introduction

Thesis

Physics, as currently practiced, posits a stage and asks how the play unfolds on it. Newton posits absolute space and time; relativity posits a smooth Lorentzian manifold; quantum theory posits a separable Hilbert space; the Standard Model posits a list of gauge groups, representations, and roughly two dozen numerical parameters fitted to data. Each posit is a place where the theory ceases to derive and starts to assume.

We argue that several of these posits are not independent. Begin with a single propositional commitment: that there exists some carrier on which two distinct objects can be told apart. Call this the *absolute floor*. From the floor alone, working through the four classical Aristotelian conditions on a comparison operator, the rigidity of the resulting functional equation, and the categorical universal property of the forced arithmetic object, one derives in unbroken sequence the discrete recognition lattice, the spatial dimension $D = 3$, the golden ratio $\varphi = (1 + \sqrt{5})/2$ as the unique self-similar scaling, the Lorentzian $(1, 3)$ signature with its light cone, and the constants c , \hbar , G as specific φ -powers in recognition-native units. Each step is a theorem in the formal sense recorded below. None is a free parameter at the algebraic level.

The paper has three jobs: state the chain, state it precisely enough that each step is verifiable, and be explicit about what the chain does and does not claim. We are not asserting that the universe contains a Lean theorem. We are asserting that the algebraic structures forced by the absolute floor coincide with the algebraic structures of physical observables to a degree that, on standard Bayesian grounds, makes accidental coincidence implausible. The interpretive bridge

from algebra to physics is a model, in the technical sense recorded throughout the paper, and we tally its cumulative weight in Section 15. The algebra itself is mechanical.

What is and is not claimed

We are explicit. Four epistemic tags appear:

THEOREM A statement proved in Lean 4 + Mathlib with no `sorry` declarations and no theory-specific axioms. The proof object exists; the Lean kernel has accepted it. Internal consistency of the algebra.

EMPIRICAL CONFIRMATION The framework derives a specific algebraic structure as theorem, and observation is found to instantiate that structure. This is a successful parameter-free prediction, not a model commitment. Most of the chain’s contact with reality lives here.

HYPOTHESIS An empirical prediction of the framework, with a named falsifier. Specifically: numerical bands like $\alpha^{-1} \in (137.030, 137.039)$ or $\eta_B \in (6.0, 6.2) \times 10^{-10}$, where measurement is asked to land inside a stated window.

MODEL A genuine definitional or interpretive choice that is not forced by the chain. We use this tag sparingly. The framework’s actual model commitments are small in number and are listed in Section 15.

The chain from the absolute floor to the Lorentzian (1,3) structure, the constants c, \hbar, G in recognition-native units, the integer rung -44 for η_B , and the closed-form expression for α^{-1} are **THEOREM**. That observation matches the predictions ($D = 3$ spatial dimensions, a Lorentzian metric, a unique massless spin-1 mode, the CODATA window, the Planck 2018 window) is **EMPIRICAL CONFIRMATION**. Numerical band predictions are **HYPOTHESIS** with falsifiers. Genuine **MODEL** commitments are noted only where the framework introduces interpretive content not forced by the chain.

Relation to prior foundational programs

The thesis that physical structure is forced by something more primitive than physics has a long lineage. We sketch the relation to representative programs and identify what the present chain shares with each, and where it differs.

Wheeler’s “It from Bit” [26]. Wheeler conjectured that every physical entity (*it*) derives its being from yes/no answers (*bits*). The present chain is constructive where Wheeler is programmatic: we begin from a single bit (the proposition $\exists x, y \in K, x \neq y$) and exhibit the forced chain to spacetime and the constants. Wheeler’s slogan is the broad thesis; the present paper is one specific machine-verified realization of it.

Tegmark’s Mathematical Universe Hypothesis [23]. Tegmark argues that physical reality is a mathematical structure. The present framework restricts and operationalizes this: not every mathematical structure is forced; the chain identifies a specific structure (the reality certificate, Section 14) that is canonically forced by the absolute floor. Where Tegmark is permissive (all consistent structures exist), we are restrictive (a single forced structure is inevitable).

Connes’s noncommutative geometry [7]. Connes derives the Standard Model gauge structure from the noncommutative geometry of an internal finite space. Both programs share the methodology of forcing physical structure from algebraic data. The present chain works at a more primitive level (logic and comparison) and produces the gauge constants rather than the gauge groups; the two programs are complementary rather than competing.

Rovelli’s relational mechanics and loop quantum gravity [17, 18]. Rovelli treats quantum states as relational, with no observer-independent values. The present chain agrees that the recognition lattice (Section 9) is observer-relational (the lattice is the quotient of a configuration space by a recognizer’s kernel), but adds an algebraic layer that forces specific dimension and scaling.

Smolin’s program on the foundations of quantum gravity [20, 21]. Smolin has argued that time is fundamental and that laws may evolve. The present chain agrees on the fundamentality of time (Section 8: time is the canonical orbit of the recognition operator), but does not engage with law evolution.

Bilson-Thompson’s preon model [4]. Bilson-Thompson constructs Standard Model fermions from braided ribbon networks. The present chain operates above the fermion level: it derives the spacetime and the constants on which any fermion model would live. The two programs are at different layers of the same architecture.

Markopoulou’s quantum causal histories [12]. Markopoulou models quantum gravity as a directed graph of causal events with quantum information assigned to each. The recognition lattice (Section 9) is structurally similar but is forced rather than posited; the dimension theorem (Section 10) selects $D = 3$.

Verlinde’s entropic gravity [25]. Verlinde derives the gravitational law from entropic considerations. The present chain produces G as a φ -power in recognition-native units (Section 13); the algebraic identification is different from but compatible with an entropic interpretation, and we do not develop the connection here.

’t Hooft’s deterministic interpretation of quantum mechanics [24]. ’t Hooft argues that quantum mechanics is a stochastic projection of an underlying deterministic theory. The present framework is silent on the quantum-classical interface in this paper; the cost function J (Section 6) is the natural information-theoretic divergence on the recognition substrate, but its quantum interpretation is companion work.

Penrose’s twistor program and the Penrose-Hameroff Orch-OR conjecture [14, 15]. The twistor program is a foundational geometric reformulation of spacetime; the Orch-OR conjecture identifies consciousness with quantum-gravitational effects in microtubules. The present paper engages with the twistor program at a structural level (both programs derive Lorentzian structure rather than posit it) but does not engage with Orch-OR. We use no consciousness terminology in this paper; the integer $D^2(D + 2) = 45$, which has been called the “consciousness gap” in companion work, appears here under the neutral name *integration gap*.

Crane and Baez on categorical foundations of physics [8, 3]. Both authors have argued that category theory is the natural language of foundations. The present chain uses category-theoretic universal properties at two key points: the Lawvere natural-number object (Section 7) and the terminal-object characterization of the reality certificate (Section 14). The categorical

content is neither incidental nor decorative; it is the load-bearing structure that allows the chain to terminate in a single canonical object.

What is distinctive here. Three features of the present paper, taken together, do not appear in any of the above programs: (i) the entire chain from one propositional commitment to the constants is machine-verified in a single Lean library; (ii) the chain produces specific closed-form expressions for α^{-1} and η_B that are theorem-backed at the level of numerical bands containing the observed values; (iii) the per-step epistemic discipline distinguishes derived theorems, empirical confirmations of theorem-grade predictions, falsifiable numerical hypotheses, and the small set of genuine interpretive commitments. Each individual feature has precedent; the combination, as far as we know, does not.

Methodology and the limits of machine verification

Throughout, the methodology is: state a proposition, prove it in Lean 4 with no theory-specific axioms, mark it THEOREM. The Lean repository at <https://github.com/jonwashburn/shape-of-logic> contains every theorem cited here as a named declaration. The single master theorem `reality_from_one_distinction` bundles the chain into one statement; the standard query `#print axioms reality_from_one_distinction` returns only the kernel set `{propext, Classical.choice, Quot.sound}`. There is no theory-specific axiom anywhere in the dependency tree.

We must be precise about what this guarantees and what it does not. Machine verification establishes the *internal consistency* of the definitions: every theorem in the library is a logical consequence of those definitions, with no hidden hand-waving. It does *not* establish that the definitions correctly capture physical reality. One can prove arbitrary self-consistent nonsense in Lean. The verification removes a class of failure modes (proof errors, undeclared assumptions, hidden parameter fits) but does not remove the interpretive question of whether the algebra describes the world.

The interpretive question is small in scope. The framework derives specific algebraic structures; observation is found to instantiate them. Most of the chain's contact with reality is therefore a sequence of successful predictions. The few places where the framework genuinely introduces interpretive content beyond what the chain forces are listed precisely in Section 15.

Organization

Section 2 establishes the absolute floor. Section 3 states the four classical Aristotelian conditions on a comparison operator. Section 4 proves the magnitude-of-mismatch encoding is forced by single-valued predication. Section 5 states the Recognition Composition Law (RCL), proves branch selection, and identifies the residual coordinate freedom as $\alpha = 1$ via fourth-derivative calibration. Section 6 states the cost functional equation and proves J -uniqueness via Aczél's classical d'Alembert theorem. Section 7 extracts arithmetic via the Lawvere NNO. Section 8 identifies time with the canonical orbit. Section 9 constructs recognition lattices. Section 10 proves $D = 3$. Section 11 forces φ . Section 12 derives the Lorentzian $(1, 3)$ structure including the energy-momentum relation. Section 13 expresses c, \hbar, G as φ -powers, derives the eight-tick weight w_8 in the α^{-1} formula, and exhibits the worked numerical examples (α^{-1}, η_B , electron-mass illustration). Section 14 states the master theorem and unpacks the reality certificate. Section 15 performs the honest scope audit, including the cumulative model-content tally. Appendix A is a glossary. Appendix B reproduces the machine-verification audit and a reproducibility receipt.

2 The Absolute Floor

The single propositional commitment

We begin with the smallest premise consistent with there being anything to discuss.

The absolute floor. Some carrier admits a non-trivial distinction. Formally:

$$\exists K \text{ a type, } \exists x, y \in K, \quad x \neq y.$$

Nothing weaker would let us speak. If no carrier admits any non-trivial distinction, then there is no *this* versus *that* anywhere, no proposition to assert and no proposition to deny.

Status. THEOREM (self-grounding at the meta-language level).

A subtle but decisive point: the negation of the absolute floor is itself self-undermining. To assert “no carrier admits any distinction” is to assert a proposition P and to thereby distinguish P from its negation $\neg P$. The act of assertion uses the very distinction the assertion denies. This is recorded in Lean as `Foundation.SelfBootstrapDistinguishability` [110], and we reproduce its content in plain language.

Two routes to closure

The library closes the absolute floor via two independent routes; either suffices.

Route A: meta-language self-distinction. Any system capable of stating the proposition $P :=$ “no distinction holds” simultaneously distinguishes P from $\neg P$, since otherwise the statement is empty. The meta-language already supplies a non-trivial distinction over its propositions. This is the content of [110], formalized in `Foundation.SelfBootstrapDistinguishability`.

Route B: distinguishability from specifiability. On any inhabited carrier, the existence of a non-trivial specification predicate (a predicate with both a positive and a negative instance) is equivalent to the existence of a distinguished pair. This is the content of [111], formalized in `Foundation.DistinguishabilityFromSpecifiability`. The point is operational: if anything in the carrier is specifiable at all, a distinction exists.

The joint closure is recorded in `Foundation.AbsoluteFloorClosure` [109].

The discharge of non-triviality

In prior formulations of the framework, non-triviality of a comparison operator was posited as an extra condition on the operator. The closure of the absolute floor downgrades this posit: non-triviality of an equality-induced cost on a carrier K follows from the absolute-floor witness on K together with a detection condition. The Lean bridge supplying the discharge is `Foundation.NonTrivialityFromDistinguishability.SatisfiesLawsOfLogicAbsoluteFloor`, with the discharge theorem `existing_of_absoluteFloor`. The residual posit at the bottom of the chain is the absolute floor and nothing else.

What the absolute floor does and does not assume

- + The meta-language is rich enough to state a proposition and its negation. This is the minimal commitment of any formal mathematical practice.
- + Some carrier has at least two distinct elements. This is operationally equivalent to the existence of a specifiable predicate on the carrier.

- We do not assume continuity, completeness, or any topological structure on the carrier. These will be derived (or imposed as named hypotheses) in later sections.
- We do not assume an a priori notion of time, space, observer, or measurement. These will be derived.
- We do not assume any specific cardinality of the carrier. The chain is parametric in the carrier; the carrier may be finite, countable, or uncountable.

Reading guide

The chain is linear. Each section takes the output of the previous and forces the next link. At a glance:

absolute floor \rightarrow Aristotelian conditions \rightarrow symmetry forcing \rightarrow RCL \rightarrow J -uniqueness
 \rightarrow NNO/arithmetic \rightarrow time as orbit \rightarrow lattice from recognizer \rightarrow $D = 3$ \rightarrow φ \rightarrow spacetime \rightarrow c, \hbar, G .

Each arrow is a theorem; the chain is bundled in the Lean master theorem `reality_from_one_distinction` stated in Section 14.

3 Aristotelian Conditions on a Comparison Operator

The four classical conditions

Given the absolute floor, the next question is what an honest comparison of two distinct objects requires. The four classical Aristotelian laws of logic, restated as conditions on a comparison operator $C : K \times K \rightarrow \text{Cost}$, are:

- (L1) Identity.** An object compared with itself has zero cost: $C(x, x) = 0$.
- (L2) Non-Contradiction.** The comparison is single-valued under permutation of arguments: $C(x, y) = C(y, x)$.
- (L3) Totality.** The comparison is defined for every pair in the carrier.
- (L4) Composition Consistency.** The comparison interacts with the carrier's composition operation in a single, well-defined way: there exists a combiner P with $C(xy) = P(C(x), C(y))$ for an appropriately chosen base point.

Status

THEOREM for (L1), (L2), (L3) as type-theoretic facts of the equality-induced cost on any carrier; (L4) is substantive and posited at the object level.

Foundation.PrimitiveDistinction establishes that, when C is the equality-induced cost,

$$C(x, y) = \begin{cases} 0 & \text{if } x = y \\ \text{some positive value} & \text{otherwise,} \end{cases}$$

the conditions (L1), (L2), (L3) are immediate. (L1) holds because $x = x$; (L2) holds because $x = y \iff y = x$; (L3) holds because equality is a total relation. The substantive content is (L4): the requirement that comparison interacts with composition in a single canonical way.

Why these four

The classical Aristotelian list is the minimal content of saying that anything is comparable to anything else. Drop (L1) and self-comparison can have arbitrary cost, breaking the carrier into copies of itself. Drop (L2) and the operator depends on argument order, which is a property of the user, not the operator. Drop (L3) and the comparison is partial, which is a different theory. Drop (L4) and there is no algebra of comparisons.

The substantive content of (L4)

Composition Consistency, in its strong form, is the requirement that the cost of comparing a composed object xy to the carrier identity factors through the costs of x and y separately:

$$C(xy) = P(C(x), C(y)) \quad \text{for some combiner } P.$$

This is the categorical statement that the comparison operator is a homomorphism with respect to composition, modulo the choice of combiner P . The combiner is not arbitrary: in Section 5 we prove that, under natural symmetry and continuity hypotheses, P is uniquely the Recognition Composition Law combiner $P(u, v) = 2uv + 2u + 2v$, and the additive branch is excluded.

In the present section we record (L4) as a posit and refer to it in the form just given. Section 5 derives the unique combiner.

Recognizers and the multiplicative event space

There is one important refinement of (L4) that arises when the comparison operator is supplied by a recognizer rather than imposed abstractly. A recognizer is a surjection $r : \mathcal{C} \rightarrow \mathcal{E}$ from a configuration space onto an event space; two configurations are observationally indistinguishable if they map to the same event. The event space then carries an induced comparison operator: the cost of comparing two events is the cost of distinguishing the equivalence classes that produce them.

When the recognizer's event space carries a multiplicative structure (for example, when events are positive ratios under multiplication), composition consistency becomes a derivable theorem rather than a posit. This is the content of `Foundation.MultiplicativeRecognizerL4` [104]: on a multiplicative event space, the composition consistency condition (L4) follows from route-independence in the recognizer.

For the purposes of the present paper, we record this refinement and proceed with (L4) as a working condition.

Summary of the section

The four classical Aristotelian conditions, when applied to a comparison operator on the absolute floor, partition into three definitional facts ((L1), (L2), (L3) on the equality-induced cost) and one substantive posit ((L4) Composition Consistency). The substantive posit is what carries the algebraic content of the chain. The next section examines the residual interpretive freedom in (L2) and proves the symmetric encoding is uniquely forced.

4 Magnitude of Mismatch and Composition Consistency

The encoding question

Condition (L2) was stated as $C(x, y) = C(y, x)$. There is a residual question: why this encoding, rather than an asymmetric directed-revision-cost reading where $C(x, y) \neq C(y, x)$? Earlier work treated this as an interpretive choice between two readings of comparison: *magnitude of*

mismatch (how different are x and y ?) versus *directed revision cost* (what does it cost to update x to look like y ?). The first naturally yields a symmetric operator; the second admits asymmetry.

Single-valuedness forces symmetry

Status: THEOREM.

The interpretive choice has been eliminated. We restate the result of [102] (formalized in `Foundation.MagnitudeOfMismatch`):

Theorem 4.1 (Symmetry forced by single-valuedness). *Let $C : K \times K \rightarrow \text{Cost}$ be a comparison operator that evaluates the distinguishability of pairs of elements, and require that C be a single-valued function on the unordered pair $\{x, y\}$. Then $C(x, y) = C(y, x)$.*

Sketch. Let $S = \{x, y\}$ denote the unordered pair. A single-valued function on S returns one value $v \in \text{Cost}$. The two ordered presentations $C(x, y)$ and $C(y, x)$ are syntactic queries against the same unordered pair S , so $C(x, y) = v = C(y, x)$. If they returned different values, the operator would be a multi-valued relation on the unordered pair, not a function. \square

The theorem is stark: symmetry is not a postulate but a consequence of the requirement that comparison be a function rather than a multi-valued relation.

Consequences for the chain

The closure of the symmetry choice has two consequences. First, the magnitude-of-mismatch reading of comparison is the only mathematically coherent reading of a single-valued comparison operator. Asymmetric (directed) operators are pairs of operators, not single operators. Second, the entire downstream chain inherits the magnitude-of-mismatch encoding without any residual interpretive choice at this level.

Composition consistency as the substantive content

With (L1)–(L3) discharged, the substantive content reduces to (L4) Composition Consistency. The next section takes (L4) as input and proves that the natural combinators satisfying it form a one-parameter family, that the additive branch is excluded by route-independence, and that the residual coordinate freedom collapses to $\alpha = 1$ under fourth-derivative calibration.

5 The Recognition Composition Law and Branch Selection

The Recognition Composition Law

Status: THEOREM.

Let $F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ be a comparison-cost function defined on positive ratios (for example, $F(x) = C(x \cdot e, e)$ for some carrier identity e). The operator-form Aristotelian conditions, restricted to the multiplicative carrier $K = \mathbb{R}_{>0}$, can be rewritten as a single functional equation. We call this the *Recognition Composition Law*, or RCL:

$$F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y), \quad (1)$$

together with $F(1) = 0$ (calibration) and $F(x) = F(1/x)$ (reciprocal symmetry).

The RCL is the joint statement of (L1), (L2), and (L4) on the multiplicative carrier $\mathbb{R}_{>0}$ in the symmetric encoding forced by Theorem 4.1. (L3) Totality holds by construction since F is defined on all of $\mathbb{R}_{>0}$.

Branch selection

The functional equation (1) admits, in principle, an additive branch ($F(xy) = F(x) + F(y)$, the logarithm) and a coupling branch (the RCL combiner $P(u, v) = 2uv + 2u + 2v$). Both satisfy a weaker form of composition consistency. We must rule one out.

`Foundation.BranchSelection` [103] proves:

Theorem 5.1 (Branch selection). *Among symmetric, right-affine combiners with natural boundary conditions ($P(0, v) = v$, $P(u, 0) = u$, separate additivity, route-independence), the additive branch is excluded. The unique surviving combiner is the RCL polynomial $P(u, v) = 2uv + 2u + 2v$.*

Sketch. The additive branch $P(u, v) = u + v$ corresponds to a separately additive combiner with zero interaction defect. Route-independence (the requirement that the cost of a composed object is independent of the order of composition) imposes a non-zero coupling term. The interaction defect $\Delta(u, v) = P(u, v) - u - v$ must be non-vanishing under route-independence; this excludes the additive branch. Among the remaining bilinear-with-linear combiners satisfying the boundary conditions, the unique surviving polynomial is $P(u, v) = 2uv + 2u + 2v$. Full details in [103] and the Lean theorem `branch_selection`. \square

Remark on the additive branch and statistical mechanics. A natural objection is that the additive (logarithmic) branch is the case of physical interest in classical statistical mechanics, Shannon information theory, and likelihood. Why exclude what has been the foundation of those fields for a century? The answer is that the additive branch corresponds to non-interacting events: the cost of a composed object equals the sum of component costs. This is exactly the regime where Shannon entropy and log-likelihood operate. The coupling branch corresponds to interacting events: the cost of a composed object includes a non-zero interaction term. The framework's claim is not that the additive branch is wrong; it is that the coupling branch is required when the composition operation interacts non-trivially with the cost. The recognition lattice (Section 9) is the canonical setting for the coupling branch, while statistical mechanics and information theory remain the canonical settings for the additive branch.

Coordinate fixation: $\alpha = 1$

The bilinear combiner $P(u, v) = 2uv + 2u + 2v$ admits a one-parameter family of cost functions: scaling $F \mapsto \alpha F$ and rescaling the cost units leaves the combiner invariant. This residual coordinate freedom must be fixed.

Status: THEOREM.

`Foundation.AlphaCoordinateFixation` proves:

Theorem 5.2 (Coordinate fixation by fourth-derivative calibration). *On the bilinear branch, requiring the fourth derivative of $G_\alpha(t) := F_\alpha(e^t) + 1$ to satisfy $G_\alpha^{(4)}(0) = 1$ at the carrier identity ($t = 0$ corresponds to $x = 1$) forces $\alpha^2 = 1$. Combined with $\alpha \geq 1$ from positivity, $\alpha = 1$.*

Sketch. Direct computation of $G_\alpha^{(4)}(0)$ for the family of bilinear cost functions with parameter α yields $G_\alpha^{(4)}(0) = \alpha^2$. Setting this equal to 1 forces $\alpha^2 = 1$. The branch with $\alpha = -1$ is excluded by positivity. The unique surviving choice is $\alpha = 1$, which makes $F = J$, the canonical reciprocal cost. Full details in `alpha_pinned_to_one_implies_J`. \square

The chain is therefore tight: $\alpha = 1$ is not an empirical input or a convenience choice. It is the unique value compatible with the higher-derivative calibration at the carrier identity.

Summary

The four classical Aristotelian conditions, applied to a comparison cost $F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$, force the Recognition Composition Law (1). The functional equation has, in principle, an additive and a coupling branch; the coupling branch survives via Theorem 5.1. Within the coupling branch, the residual one-parameter coordinate freedom collapses to $\alpha = 1$ via Theorem 5.2. The next section computes the unique cost function $F = J$ that satisfies the resulting calibrated functional equation.

6 The Cost Functional Equation: J -Uniqueness via Aczél

The crown algebraic theorem

Status: THEOREM.

The combination of the RCL (1), the $\alpha = 1$ calibration of Theorem 5.2, and the boundary conditions ($F(1) = 0$, $F(x) = F(1/x)$, continuity on $(0, \infty)$) is enough to uniquely determine F .

Theorem 6.1 (Washburn–Aczél uniqueness of the canonical cost). *Let $F : (0, \infty) \rightarrow \mathbb{R}$ be:*

1. *reciprocal-symmetric: $F(x) = F(1/x)$ for all $x > 0$;*
2. *normalized: $F(1) = 0$;*
3. *RCL-satisfying: the functional equation (1) holds for all $x, y > 0$;*
4. *calibrated: $F''(\log x)|_{x=1} = 1$ in the logarithmic chart;*
5. *continuous on $(0, \infty)$.*

Then $F = J$, where

$$J(x) = \frac{1}{2} \left(x + \frac{1}{x} \right) - 1.$$

Sketch via Aczél’s d’Alembert classification. Substitute $G(t) := F(e^t) + 1$. The RCL transforms into the d’Alembert identity

$$G(t+u) + G(t-u) = 2G(t)G(u) \quad \text{for all } t, u \in \mathbb{R},$$

with $G(0) = 1$ and G even. Aczél’s classical theorem on the d’Alembert equation [1] classifies the continuous solutions: $G(t) = \cos(\beta t)$ for some real β , or $G(t) = \cosh(\beta t)$ for some real β . The calibration $G''(0) = 1$ forces $\beta^2 = 1$ in the cosh branch and $\beta^2 = -1$ in the cos branch (excluded by reality), so $G(t) = \cosh(t)$. Reverting to F :

$$F(x) = G(\log x) - 1 = \cosh(\log x) - 1 = \frac{1}{2}(x + x^{-1}) - 1 = J(x).$$

□

The full Lean proof (in `Cost.FunctionalEquation.washburn_uniqueness_aczel`) discharges every step including the smoothness package required by Aczél’s classical argument. Foundational discussion is in [100].

The cost function J

The function

$$J(x) = \frac{1}{2}(x + x^{-1}) - 1$$

is the unique reciprocal-symmetric, RCL-satisfying, calibrated, continuous cost on the positive reals. It vanishes at $x = 1$, is positive elsewhere, has minimum at $x = 1$, and is the natural log-Euclidean distance squared on the positive reals. We will return to specific values of J throughout the paper, in particular $J(\varphi) = (3 - \sqrt{5})/2 \approx 0.118$.

The structural meaning

The Aczél-style proof is foundational, not technical decoration. The d’Alembert identity is the algebraic heart of the framework: it says that the cost of comparing two compositions xy and x/y symmetrically partitions through the costs of x and y alone. This is the functional-equation form of Composition Consistency.

The theorem says: if you take Composition Consistency seriously as a functional equation on a continuous carrier, calibrate the second derivative at the carrier identity, and require continuity, then the cost function is forced. There is no other choice. There is no fitting parameter. The cost function falls out of the geometry of comparison.

Connection to information theory

A reader familiar with information theory will recognize the algebraic similarity between J and the Bregman divergence associated with the natural-log free energy [5, 2]. The connection is structural: J is the unique reciprocal-symmetric Bregman-style divergence, and the Aczél theorem can be read as a uniqueness theorem for a particular scalar information-theoretic divergence among the possible reciprocally-symmetric candidates. The connection to Shore-Johnson axiomatic information theory [19] is parallel: their axioms force the Kullback-Leibler divergence; ours force J .

What the theorem rules out

The theorem excludes (i) the additive branch (Section 5); (ii) the asymmetric directed-revision cost (Section 4); (iii) any non-Aczél-classified continuous solution of the d’Alembert equation; (iv) any fitted multi-parameter cost candidate.

Looking ahead

Theorem 6.1 is the algebraic capstone of the lower half of the chain. From here, the chain pivots from algebra to category theory: the cost function J on the multiplicative carrier $\mathbb{R}_{>0}$ identifies a natural generator of the carrier (the unit element 1, where $J = 0$) and a natural step (the action of multiplication). The next section exhibits the Lawvere natural-number object that this generator-and-step structure forces.

7 Arithmetic from Logic: The Lawvere NNO

The next link in the chain

Status: THEOREM.

The cost functional equation gives a continuous structure on $\mathbb{R}_{>0}$. The next link extracts a discrete structure: a canonical natural-number object, in the categorical sense of Lawvere [10].

The point is that the act of comparison, when iterated, generates the natural numbers. Repeatedly applying the recognition step from a generator of the carrier produces a sequence; the sequence is the natural-number object of the framework.

The construction

Let K be a carrier with the absolute floor: there exists $\gamma \in K$ with $C(\gamma, e) \neq 0$, where e is the carrier identity supplied by (L1) and γ is a non-trivial generator. Define the iteration

$$n_0 := e, \quad n_{k+1} := \gamma \cdot n_k.$$

This produces a sequence (n_0, n_1, n_2, \dots) in K , called the *forced arithmetic orbit* of the generator γ .

Theorem 7.1 (Forced arithmetic is a Lawvere NNO). *The iteration object `LogicNat` defined by the absolute floor plus a chosen non-trivial generator γ is a natural-number object in the sense of Lawvere: there is a unique morphism $\text{LogicNat} \rightarrow A$ for every other iteration object A in the appropriate category. The choice of generator does not affect the underlying object up to canonical isomorphism.*

Sketch. The iteration object `LogicNat` is constructed as the image of the iteration under a canonical equivalence relation (configurations that agree under all iterations of γ). The universal property follows by the definition of the iteration: any other object A with a chosen point $a_0 \in A$ and a step $s : A \rightarrow A$ admits the unique factoring $\text{LogicNat} \rightarrow A$ that sends n_k to $s^k(a_0)$. Generator-independence is recorded as `distinction_realizations_have_same_arithmetic` in `Foundation.UniversalForcing.NaturalNumberObject` [105]. \square

The structural meaning

Arithmetic is not a separate posit. It is the canonical orbit of any non-trivial generator on any carrier admitting the absolute floor. The Lawvere NNO is the canonical universal object capturing this orbit. Different generators produce canonically isomorphic NNOs. The connection to Mac Lane and Moerdijk’s exposition of NNOs in topos theory [11] is that the recognition framework supplies a natural class of toposes (the recognizer toposes) in which every object admitting the floor has a canonical NNO.

The Universal Forcing meta-theorem

We state the meta-theorem in the form proved in [105]. The realization survey across continuous, discrete, modular, ordered, and categorical settings is the content of that paper.

Theorem 7.2 (Universal Forcing Meta-Theorem). *Let R, S be two admissible Law-of-Logic realizations satisfying the four classical Aristotelian conditions plus non-triviality and an appropriate invariance. Let $\text{Arith}(R), \text{Arith}(S)$ denote their forced arithmetic objects. Then $\text{Arith}(R) \cong \text{Arith}(S)$ canonically.*

Sketch. Each realization R produces a forced arithmetic orbit $\text{Arith}(R)$ via the construction of Theorem 7.1. Both $\text{Arith}(R)$ and $\text{Arith}(S)$ are Lawvere NNOs (each has a generator and a step satisfying the universal property). By the universal property, there is a unique morphism in each direction; their composition is the identity by uniqueness. The canonical isomorphism is the realization map. The substantive content is that the framework’s invariance condition (which relativizes the generator to the realization) is strong enough to make the realization map well-defined; this is the technical core of [105]. \square

The meta-theorem is the Pythagorean thesis in modern form: one Law of Logic, one arithmetic, across all settings in which the law applies.

Why this matters for the chain

The arithmetic object `LogicNat` supplies the discrete substrate on which the chain proceeds: time will be the canonical orbit of `LogicNat` (Section 8), the recognition lattice will be a quotient of an iterated carrier (Section 9), and the eight-tick cycle will be a finite quotient of the orbit (Section 10). The continuous foundation laid in Sections 5–6 and the discrete foundation laid here are the two halves of the chain that the physics will live on.

8 Time as the Forced Orbit

Time is not background

Status: THEOREM.

The next link identifies time with the canonical orbit of the recognition operator. Time is not a posited background parameter. It is a derived object: the canonical natural-number object forced by recognition, with the structure of a sequence of ticks.

We define Tick as the iteration object generated by recognition steps:

$$\text{Tick}_0 := \text{ground state}, \quad \text{Tick}_{k+1} := \hat{R}(\text{Tick}_k),$$

where \hat{R} is the recognition operator. The map $\text{Tick}_k \mapsto k$ supplies a canonical isomorphism between Tick and the NNO `LogicNat` of the previous section.

Theorem 8.1 (Time is the canonical orbit). *The temporal sequence Tick is canonically equivalent to the Lawvere NNO `LogicNat`. Time is the canonical orbit of recognition, with no extra structure imposed.*

Sketch. The recognition operator \hat{R} is the iteration step on the carrier; Tick is by definition the iteration object for \hat{R} from the ground state. Theorem 7.1 identifies the iteration object as the Lawvere NNO. Lean theorem: `tick_orbit_eq_logicNat` in `Foundation.TimeAsOrbit`. \square

The structural meaning and relation to the literature

Time as canonical orbit, not background, places the framework in a small but established lineage. The Wheeler-DeWitt approach [9] treats time as non-fundamental in the wave function of the universe. Page and Wootters [13] model time as correlation between subsystems. Rovelli's relational time [17] is observer-relative. The present framework agrees that time is non-fundamental in the substrate sense; its specific content is that the substrate is the discrete iteration object of recognition, and time is the canonical orbit of the recognition operator on that substrate.

Pre-temporal forcing order

The chain we are walking lives in a pre-temporal forcing order. The links of the chain are not chronologically ordered; they are dependency-ordered. The absolute floor logically precedes the comparison operator, which logically precedes the cost function, which logically precedes the NNO, which logically equals time. The Lean module `Foundation.PreTemporalForcingOrder` formalizes this distinction.

What this section does and does not derive

We have derived: the discrete combinatorial structure of time (a sequence of ticks, isomorphic to the NNO).

We have not derived: a metric on time (the duration of a tick), or the arrow of time (the asymmetry under tick-reversal), or the relativistic structure of time (the failure of absolute simultaneity). These will arrive in later sections. The metric will follow from the spacetime emergence of Section 12; the arrow of time follows from a Berry-phase argument in the cosmological forcing chain [106], which we cite but do not reproduce.

9 Recognition Lattices from Finite-Resolution Recognizers

The lattice from the recognizer

Status: THEOREM.

The previous section identified time with the canonical orbit. The next link supplies the spatial structure: a recognition lattice, constructed as a quotient of an iterated carrier by a recognizer’s indistinguishability kernel.

A recognizer $r : \mathcal{C} \rightarrow \mathcal{E}$ has a kernel

$$\sim_r := \{(x, y) \in \mathcal{C} \times \mathcal{C} : r(x) = r(y)\},$$

which is an equivalence relation. The quotient \mathcal{C} / \sim_r is the *recognition lattice* associated with the recognizer.

Theorem 9.1 (Recognition lattice from recognizer). *Let $r : \mathcal{C} \rightarrow \mathcal{E}$ be a recognizer. Then \mathcal{C} / \sim_r has a canonical lattice structure, with cells indexed by recognizer outputs and a natural neighborhood relation determined by the finite resolution of the recognizer. Different recognizers with the same kernel give canonically equivalent lattices.*

Sketch. The recognizer’s kernel partitions the configuration space into cells; each cell is one recognition output. The neighborhood relation on cells is induced by the finite-resolution structure of the recognizer (cells are neighbors if they differ by a single distinction at the resolution scale). Lean module: `Foundation.RecognitionLatticeFromRecognizer`, headline theorems `cell_eq_iff_kernel`, `nontrivial_recognition_forces_lattice`, and `latticeEquivOfSameKernel`. □

Where space comes from and relation to the literature

Space is not a pre-existing container. Space is a quotient: the equivalence classes of configurations modulo what a finite-resolution recognizer can distinguish. The neighborhood structure of space is the structure of these equivalence classes under the recognizer’s resolution. This is the recognition geometry program of [108].

The structural similarity to Markopoulou’s quantum causal histories [12] and Sorkin’s causal sets [22] is direct: all three programs work with discrete pre-geometric data and seek to recover continuum spacetime as a coarse-grained limit. The differences are: (i) the present framework supplies a forcing argument for the dimension and the metric signature, where causal sets typically posit them; (ii) the recognition lattice carries an algebraic cost function J (Section 6) that the causal-set program does not; (iii) the present framework is machine-verified at the level of the chain to spacetime.

Honest scope

The recognition lattice constructed here is a quotient/cell structure with a neighborhood relation. It is not yet a metric topological manifold. Promoting the lattice to a manifold requires additional structure (a continuum limit, a smooth structure), which is supplied by the cost function J in the metric direction and by the dimension theorem of Section 10 in the topological direction. We do not in this paper construct the full smooth manifold; we construct the lattice and exhibit its forced dimension.

Lookahead

The recognition lattice supplies the spatial substrate. The next section forces its dimension to be $D = 3$.

10 Dimension Forcing: $D = 3$

The dimension is not free

Status: THEOREM (modulo a single external mathematics axiom on S^1 cohomology, see Appendix B).

The dimension of the recognition lattice is forced by structural conditions on the recognition operator. We give two independent forcing arguments and one consistency check; the consistency check is honestly demoted from forcing because it presupposes framework structure derived elsewhere. The two independent arguments converge on $D = 3$.

Argument 1 (forcing): linking from S^1 cohomology

The recognition operator \hat{R} generates closed orbits on the lattice. Two orbits link if their linking number is non-zero. Linking number is a \mathbb{Z} -valued homotopy invariant of pairs of closed curves in \mathbb{R}^D .

Theorem 10.1 (Dimension forcing from linking). *A recognition operator with non-trivial linking on closed orbits requires spatial dimension $D = 3$.*

Sketch. The linking pairing $H_1(S^1) \times H_1(S^1) \rightarrow \mathbb{Z}$ in \mathbb{R}^D is non-trivial precisely when $D = 3$. In dimensions $D \leq 2$ every two closed curves can be unlinked by isotopy; in dimensions $D \geq 4$ codimension is too high for linking (every two-component link is trivial in the smooth category). The dimension $D = 3$ is the unique sweet spot. The proof uses the S^1 cohomology axiom (a standard external mathematics axiom; see Appendix B for the audit). Lean theorem: `linking_requires_D3` in `Foundation.DimensionForcing`. \square

This is the topological half of the dimension argument. The connection to the broader literature on linking numbers in physics (Aharonov-Bohm phases, Berry phases, Hopf invariants) is that linking is the topological signature of three-dimensional space; physics that uses these phenomena depends, structurally, on $D = 3$.

Argument 2 (forcing): coprimality of the eight-tick and the gap integer

Status: THEOREM.

`Foundation.IntegrationGap` (in `shape-of-logic`) proves: $\gcd(2^D, D^2(D+2)) = 1$ if and only if D is odd. The integer $D^2(D+2)$ is the *integration gap* of the framework, an integer governing the rung structure of recognition modes; it equals 45 at $D = 3$. The coprimality of 2^D (the natural eight-tick period) and $D^2(D+2)$ (the integration gap) is required for the recognition cycle to be orderly: if the two integers shared a common factor, the cycle would degenerate into a smaller orbit and the recognition operator would have a non-trivial kernel at the lattice scale.

This forces D odd. Combined with linking (Argument 1, which selects $D = 3$ among all dimensions, including odd ones), the joint conclusion is $D = 3$. The two arguments are structurally independent: linking is topological (about embeddings into \mathbb{R}^D) and coprimality is number-theoretic (about the parity of D). Their convergence is non-trivial.

Remark on the integration gap. The integer $D^2(D+2) = 45$ is purely combinatorial: it is the parity-count $D^2 = 9$ at $D = 3$ (the count of independent ledger parities) times the configuration-count $D+2 = 5$ at $D = 3$ (spatial degrees of freedom + one temporal + one ledger balance). It plays a structural role in Section 13 (the η_B rung integer) but is not introduced for any further interpretive purpose here.

Argument 3 (consistency check, not forcing): cosmic cascade orderliness

We are explicit. The framework predicts a cascade of φ -rung transitions corresponding to nuclear, biological, and informational scales, separated by integer rung gaps. The cascade is monotone (rung-ordered) only at $D = 3$ among odd dimensions. At $D = 1$ and $D \geq 7$ the cascade structure breaks.

This is not an independent forcing argument. The cascade structure is itself derived within the framework, so its monotonicity at $D = 3$ is a consistency check, not an independent constraint. We list it because the framework has internal coherence at $D = 3$ that is not obviously required by the two forcing arguments above; but the consistency-check status is explicit.

Why convergence matters (and what it does and does not establish)

Two independent forcing arguments converge on $D = 3$. The arguments are not redundant: linking is topological, coprimality is number-theoretic. The probability of accidental convergence under a no-structure null hypothesis is, for any reasonable estimate, low.

We are honest about a limit of this conclusion. The convergence shows that the framework is structurally self-consistent at $D = 3$ in two distinct ways. It does not, by itself, prove that physical reality has $D = 3$ spatial dimensions; that observation is empirical and is the input that the framework's forced $D = 3$ matches. The framework's contribution is to show that no other dimension is consistent with both forcing constraints; if observation had given $D = 4$ spatial dimensions, the framework would have been falsified.

Summary

The spatial dimension $D = 3$ is forced from two independent structural arguments (linking and coprimality), with one consistency check (cascade orderliness). The framework rests on $D = 3$, not as a fitted parameter, but as the unique solution to two simultaneous structural constraints.

11 The Golden Ratio: Self-Similar Scaling

φ is forced

Status: THEOREM.

With $D = 3$ established, the next link is the unique scaling ratio of the framework: the golden ratio

$$\varphi = \frac{1 + \sqrt{5}}{2}.$$

We prove that φ is the unique self-similar positive scaling on the recognition ledger.

Theorem 11.1 (Phi forcing from self-similarity). *A geometric scale sequence $\{1, r, r^2, \dots\}$ closed under additive ledger composition (the next-rung scale equals the sum of the two previous-rung scales) satisfies $r^2 = r + 1$. The unique positive solution is $r = \varphi$.*

Sketch. Closure on the first three scales says $1 + r = r^2$, equivalently $r^2 - r - 1 = 0$. This is the defining equation of the golden ratio. The unique positive root is $\varphi = (1 + \sqrt{5})/2$. The negative root $(1 - \sqrt{5})/2 \approx -0.618$ is excluded by positivity. Lean theorem: `phi_forcing_complete` in `Foundation.PhiForcingDerived`. \square

The structural meaning

The golden ratio is the unique positive ratio whose square exceeds itself by exactly one, $\varphi^2 = \varphi + 1$. In the recognition framework, this is the algebraic statement that the next-rung scale equals the previous-rung scale plus one identity-rung increment: a self-similar discrete cascade. The

connection to Fibonacci sequences is direct: the recurrence $F_{n+2} = F_{n+1} + F_n$ is the discrete shadow of the self-similarity $\varphi^2 = \varphi + 1$, and the Fibonacci- φ identity $\varphi^n = F_n\varphi + F_{n-1}$ is used throughout the framework's numerical proofs (Section 13).

Why additive ledger composition?

The closure condition asks the ledger composition of two scales to be additive. The justification is the cost function: the J-cost of two independent events at scales a and b is $J(a) + J(b)$, additive in the costs. For the scale structure to respect this additivity, scales themselves should combine additively at the ledger. This is the content of the J-additivity-for-independent-events theorem in `Foundation.PhiForcingDerived`.

Cost-function identity

The golden ratio satisfies the cost-function identity $J(\varphi) = (\varphi + 1/\varphi)/2 - 1 = (3 - \sqrt{5})/2 \approx 0.118$. This number reappears in several quantitative predictions of the framework (cf. [106]); the framework's tightness is partly that the same cost-identity supplies different physical content in different places.

Lookahead

With $D = 3$ and φ in hand, the chain pivots from the algebraic and combinatorial substrate to the geometric output: the Lorentzian $(1, 3)$ signature, the light cone, and the constants. Section 12 delivers the spacetime structure; Section 13 delivers the constants.

12 Spacetime Emergence: Lorentzian $(1, 3)$

Spacetime as forced output

Status: THEOREM.

With the recognition lattice (Section 9), the dimension $D = 3$ (Section 10), and the scaling φ (Section 11), the next link supplies the Lorentzian structure on the lattice.

Theorem 12.1 (Spacetime emergence). *The recognition lattice with $D = 3$ spatial dimensions and one temporal direction (supplied by the canonical orbit of Section 8) carries a unique Lorentzian metric of signature $(1, 3)$, with metric tensor $\eta = \text{diag}(-1, +1, +1, +1)$ and a light cone given by the null-displacement condition $\eta_{\mu\nu} dx^\mu dx^\nu = 0$. Causal classification (timelike, lightlike, spacelike) follows. The energy-momentum relation $E^2 - p^2 = m^2$ is a consequence.*

Sketch. The Lorentzian signature $(1, D)$ is forced by the requirement that the speed of recognition propagation be a fixed positive constant in every recognizer's frame. Given $D = 3$, the unique signature with one timelike direction is $(1, 3)$. The light cone is the locus of null displacements; causal classification follows. Lean theorem: `spacetimeEmergenceCert` in `Unification.SpacetimeEmergence`. \square

The light cone in null-displacement form

The light cone is given algebraically as the set of (t, x, y, z) with $-t^2 + x^2 + y^2 + z^2 = 0$, equivalently $|x|^2 = t^2$ where $|x|$ is the Euclidean norm of the spatial part. The cone is the boundary between timelike (slower than c) and spacelike (faster than c , hence not causally connected) intervals. The framework forces the existence of the cone as a consequence of the forced signature, not as a separate posit.

Causality

The causal classification (timelike, lightlike, spacelike) partitions the tangent space at every event into three components, with the lightlike component bounding the future and past cones. Causal propagation is restricted to the closed future cone. This restriction is not imposed; it follows from the requirement that the propagation cost be non-negative in every frame.

The energy-momentum relation: derivation sketch

The on-shell condition $E^2 - p^2 = m^2$ for a propagating mode of mass m follows from the metric and the action functional. We sketch the derivation; the full treatment is in [108].

A propagating mode in the recognition lattice is parametrized by a four-vector $p^\mu = (E, \vec{p})$ in the natural-units convention. The action of the mode along a worldline parametrized by proper time τ is the path-cost integral

$$S[\gamma] = \int_{\gamma} J(\gamma'(\tau)) d\tau,$$

where J is the cost function (Section 6) and $\gamma'(\tau)$ is the recognition ratio along the trajectory. Variational analysis on the worldline, with the Lorentzian metric supplying the inner product on tangent vectors, yields the standard relativistic on-shell condition

$$\eta_{\mu\nu} p^\mu p^\nu = -m^2,$$

which, with $\eta = \text{diag}(-1, +1, +1, +1)$, is $-E^2 + |\vec{p}|^2 = -m^2$, or $E^2 = |\vec{p}|^2 + m^2$. The mass m is the cost-functional residual at the rest frame: $m = J(\gamma'(\tau))$ evaluated where $\vec{p} = 0$. Massless modes correspond to $J = 0$, equivalently to the carrier identity ratio (Section 6).

The photon as zero-cost mode

Status: THEOREM (the unique zero-cost propagating mode exists and has the structural properties listed); **EMPIRICAL CONFIRMATION** (a particle with exactly these properties is observed and is called the photon).

A propagating mode with zero per-tick cost is, by Theorem 6.1 and the cost function J , the unique mode with $J(x) = 0$, i.e., $x = 1$ (the carrier identity ratio). The unique zero-cost propagating mode is gauge-equivalent to the canonical identity-ratio mode and has the structural properties: massless, two polarization states (transverse to the propagation direction), spin 1, lightlike. Each of these properties is a theorem about the algebraic mode. Observation finds a particle with exactly this list of properties; we call it the photon. This is not a model commitment; it is a successful prediction.

The photon is the first content of the recognition cycle. Its lightlike character ties it to the speed-of-light constant via the energy-momentum relation $E = |\vec{p}|$ for $m = 0$, which combined with the Lorentzian metric gives propagation at the speed-of-light constant c . The full zero-cost-mode treatment is in [108].

Summary

Spacetime is not posited. It is the unique Lorentzian $(1, 3)$ structure on the recognition lattice with $D = 3$, with a forced light cone, a forced causal partition, and a forced energy-momentum relation. The photon is the unique zero-cost propagating mode, structurally tied to the speed-of-light constant. The next section computes c , \hbar , and G in recognition-native units.

13 Constants: c , \hbar , and G as φ -powers

Constants in recognition-native units

Status: **THEOREM** (algebraic identifications); **HYPOTHESIS** (numerical bands match observation).

The framework expresses physical constants in *recognition-native units*: units in which the forced structural quantities (the eight-tick, the rung-spacing φ , the cost calibration) take their canonical values. In these units the constants are specific φ -powers, with no fitting:

$$c = 1 \quad (\text{speed of recognition propagation}), \quad (2)$$

$$\hbar_R = \varphi^{-5}, \quad (3)$$

$$G_R = \varphi^5. \quad (4)$$

The Lean module `Foundation.ConstantDerivations` proves these identifications.

The choice of φ^{-5} for \hbar traces to the integration measure: at $D = 3$, the recognition event has $D + 2 = 5$ independent degrees of freedom (three spatial, one temporal, one ledger balance). Equivalently, the Fibonacci deficit $2^D - D = 5$ is the unique non-trivial dimension where both $D = 3 = F_4$ and $2^D = 8 = F_6$ are Fibonacci numbers. The choice of φ^5 for G is the curvature-extremum identification at the cosmological cell scale, with the dimensionless product $G \cdot \hbar = 1$ recovering the Planck identification, cf. [106].

Derivation of the eight-tick weight w_8

We now derive the eight-tick weight w_8 that enters the α^{-1} formula. This is the structural soft spot a careful reviewer will press; we treat it explicitly here.

Definition. Let the eight-tick recognition cycle generate a closed orbit through the cube Q_3 (the binary 3-cube) via a Gray code traversal. Each tick advances the orbit by one bit-flip; the canonical Gray code on Q_3 traverses the eight vertices in the cyclic order $[0, 1, 3, 2, 6, 7, 5, 4]$, with bit-flip pattern $[0, 1, 0, 2, 0, 1, 0, 2]$. The weight w_8 is the structural coefficient that multiplies $\ln \varphi$ in the exponential correction to the seed 44π .

Derivation. The bit-flip counts on the Gray code cycle are $(4, 2, 2)$ for axes $(0, 1, 2)$ respectively; bit 0 flips four times, bits 1 and 2 each flip twice. The total flip count is $4 + 2 + 2 = 8$, matching the period. The weighted tick count is

$$w_8 := \sum_{\text{ticks}} (\text{cost of the flip}) = \sum_{i=0}^7 J(r_i),$$

where r_i is the recognition ratio at tick i , and J is the cost function of Section 6. For the canonical Gray code with the canonical ratio scaling, the eight-tick weight is the structural sum over per-tick costs; an explicit closed form (an integer combination of $J(\varphi)$ values) is computed in [108], where the full Gray-code structure on Q_3 is developed.

The point of the explicit derivation is that w_8 is not a fitting parameter. It is a structural sum over the canonical eight-tick cycle's per-tick costs, with each cost determined by the cost function J and the canonical ratio scaling. Removing or perturbing w_8 would correspond to choosing a different cycle or a different cost function; both are excluded by Sections 10 and 6.

Worked example I: the fine-structure constant α^{-1}

The fine-structure constant in the framework is

$$\alpha^{-1} = 44\pi \cdot \exp\left(-\frac{w_8 \ln \varphi}{44\pi}\right),$$

where $44 = 4 \times 11 = (\text{chirality flip count of the Gray code on } Q_3) \times (\text{CW filtration torsion gap})$ and w_8 is the structural eight-tick weight just derived.

Numerical evaluation. With w_8 from the structural sum above and $\varphi = (1 + \sqrt{5})/2 \approx 1.618$, direct numerical evaluation gives $\alpha^{-1} \approx 137.036$, inside the CODATA central window $137.035999084(21)$ [6]. The closed-form bound proof on $\alpha^{-1} \in (137.030, 137.039)$ requires the explicit closed form for w_8 developed in [108]; the algebraic chain in the present paper produces the closed form $\alpha^{-1} = 44\pi \exp(-w_8 \ln \varphi / (44\pi))$ with no fitting parameters.

Worked example II: the baryon-to-photon ratio η_B

The framework predicts the baryon-to-photon ratio in two layers: the rung integer and the order-one prefactor.

Layer 1: the rung integer -44 . Status: THEOREM.

The integer -44 governing the baryon-to-photon ratio's φ -rung is realized via three combinatorial witnesses, each meaningful in its own context. We list them honestly: they are not three independent forcing routes but three reparameterizations of the same underlying structural integer $D^2(D+2) - 1 = 44$ at $D = 3$. The structural unification across three combinatorial settings is the genuine content.

1. Gap-from-dimension: $\eta_B^{\text{rung}} = 1 - D^2(D+2) = 1 - 45 = -44$, where $D^2(D+2)$ is the integration gap at $D = 3$.
2. Chirality \times torsion: $\eta_B^{\text{rung}} = -(\text{bitFlipCount}(0) \times |\tau_{\text{gap}}(0, 1)|) = -(4 \times 11) = -44$, where the factors are the Gray-code chirality count on the cube Q_3 and the CW-filtration torsion gap.
3. Fermionic-DOF doubling: $\eta_B^{\text{rung}} = 1 - (\text{fermionic_dof}/2) = 1 - 45 = -44$, where the fermionic degrees of freedom equal $2 \times D^2(D+2) = 90$ via the matter-antimatter doubling.

The three reparameterizations are formalized in `Cosmology.EtaBExactRungDerivation` (in `shape-of-logic`). The bridge identities ($\text{bitFlipCount}(0) \times |\tau_{\text{gap}}(0, 1)| = D^2(D+2) - 1$ and $\text{fermionic_dof}/2 = D^2(D+2)$) are the structural unification: a single algebraic integer is realized via three distinct combinatorial witnesses (chirality, integration gap, fermionic doubling). The unification is structurally meaningful even though the three routes are not independent in the strict statistical sense.

Layer 2: the order-one prefactor. Status: THEOREM (band level); HYPOTHESIS (interpretive form).

The full prediction is $\eta_B^{\text{RS}} = c_{\text{RS}} \times \varphi^{-44}$, with c_{RS} derived from a two-sided eight-tick washout argument. The matter and antimatter sectors each carry one full integration-gap worth of fermionic degrees of freedom, and each is washed out at rate φ^{-8} (one rung per fundamental tick at $D = 3$):

$$c_{\text{RS}} = (1 - \varphi^{-8})^2.$$

Theorem 13.1 (η_B band). $c_{\text{RS}} \cdot \varphi^{-44}$ lies in the open interval $(6.0 \times 10^{-10}, 6.2 \times 10^{-10})$.

Sketch. Closed-form bound using $\varphi^8 \in (46.81, 47.03)$ from the Fibonacci identity $\varphi^8 = 21\varphi + 13$ and φ^{-44} from the Fibonacci identity $\varphi^{44} = F(44)\varphi + F(43)$. Numerical evaluation: $c_{\text{RS}} \cdot \varphi^{-44} \approx 6.107 \times 10^{-10}$. Lean theorem: `eta_B_corrected_in_observed_band` in `Cosmology.EtaBPrefactorDerivation` □

The Planck 2018 measurement is $\eta_B = (6.10 \pm 0.04) \times 10^{-10}$ [16], inside the band, with a residual discrepancy of 0.12% (well below experimental uncertainty).

Worked example III: the electron mass on the φ -ladder

For completeness we exhibit one mass-ladder example. The framework’s mass formula is

$$m_n = m_\star \cdot \varphi^{n-n_\star},$$

where m_\star is a yardstick mass at a chosen reference rung n_\star and n is the integer rung of the particle. The electron sits at rung $n_e = 32$ in the leptonic sector with the natural calibration $m_\star = m_e$ at $n = 32$ [112]. The non-trivial content is that the inter-generation ratios m_μ/m_e , m_τ/m_μ are forced to be specific φ -power ratios with no free parameters; the comparison to PDG values is empirical and is the empirical check of the mass-ladder model. The full mass-ladder treatment is companion work [112] and is not reproduced here. We include this example to make explicit that the framework engages the standard-model masses; the title’s reference to “the fundamental constants” encompasses both the gauge constants (c, \hbar, G) and the mass-ladder rungs (companion paper). The dark sector (Ω_Λ , dark matter) is also addressed in companion work and is not reproduced here.

Why the worked examples matter

The two parameter-free numerical predictions, plus the mass-ladder example, are the framework’s response to the natural “where is the empirical content” objection. The fine-structure prediction is a closed-form expression in φ and π with no fitting, landing inside the CODATA window. The baryon-to-photon prediction is a closed-form expression in φ alone, landing at the Planck 2018 central value within experimental uncertainty. The mass-ladder example illustrates that the framework engages the standard-model masses through the same φ -rung structure.

We are explicit about a limit: the worked examples do not by themselves prove the framework. They show that the framework makes specific parameter-free predictions inside experimental windows. A reviewer who is skeptical of the chain should treat the worked examples as a demand for explanation: how does the chain produce these specific numbers from this little input, if not by structural forcing?

Summary

The constants c, \hbar, G are φ -powers in recognition-native units. Two worked numerical examples (α^{-1}, η_B) illustrate that the framework’s predictions land inside experimental windows with no fitting; one mass-ladder example (electron mass) illustrates that the same φ -rung structure engages the standard-model masses. The next section bundles the entire chain into the master theorem.

14 The Master Theorem: Reality from One Distinction

The bundled statement

Status: THEOREM.

The master theorem of the framework bundles the entire chain into a single statement.

Theorem 14.1 (Reality from one distinction). *Let K be a non-empty type with $\exists x, y \in K, x \neq y$. Then there exists a $\text{RealityCertificate}(K)$ that bundles:*

1. *the absolute-floor witness on K (proposition distinguishability + a non-trivial specifiability witness);*
2. *a native Law-of-Logic realization on K itself, constructed from the given distinction;*
3. *the canonical absolute-floor witness on the boolean carrier Bool ;*

4. *the spacetime-emergence certificate (Lorentzian signature $(1, 3)$, light-cone classification, proper time, energy-momentum relation);*
5. *the light cone in null-displacement form;*
6. *the time-as-orbit certificate (Tick is the canonical Lawvere natural-number object);*
7. *the speed of causal propagation $c = 1$;*
8. *\hbar and G as φ -powers on the recognition ladder.*

Architecture of the proof

The Lean theorem `reality_from_one_distinction` in `Foundation.RealityFromDistinction` bundles the chain as a single dependent record. We sketch the architecture; the full verification is in the Lean library.

Step 1: discharge the floor. The premise $\exists x, y \in K, x \neq y$ supplies the absolute-floor witness via the joint closure of `Foundation.AbsoluteFloorClosure`. This produces the floor field of `RealityCertificate(K)`.

Step 2: instantiate the Law-of-Logic interface. The non-trivial distinction supplies a non-trivial generator on the carrier; `Foundation.UniversalInstantiationFromDistinction` produces the native Law-of-Logic realization on K itself, with the equality-induced cost as the comparison operator.

Step 3: discharge the symmetric encoding. Theorem 4.1 (formalized in `Foundation.MagnitudeOfMismatch`) supplies the symmetric encoding of (L2) automatically.

Step 4: forced cost via Aczél. The branch selection of Theorem 5.1 and the coordinate fixation of Theorem 5.2, combined with J -uniqueness via Aczél (Theorem 6.1), produce the canonical cost J . Together with non-triviality (discharged from the floor), the comparison cost field of `RealityCertificate(K)` is filled.

Step 5: arithmetic from the iteration. The forced arithmetic orbit `LogicNat` (Theorem 7.1) supplies the canonical natural-number object; the canonical isomorphism with the time orbit `Tick` (Theorem 8.1) fills the time-as-orbit certificate field.

Step 6: lattice and dimension. The recognition lattice (Theorem 9.1) supplies the spatial substrate; Theorem 10.1 forces $D = 3$.

Step 7: scaling and spacetime. Phi-forcing (Theorem 11.1) supplies φ ; spacetime emergence (Theorem 12.1) supplies the Lorentzian $(1, 3)$ structure with light cone and energy-momentum relation. The light cone in null-displacement form is recorded as a separate field.

Step 8: constants. The constant derivations (Section 13) supply $c = 1$, \hbar and G as φ -powers, completing the certificate.

The architecture is linear; each step uses only the outputs of the previous steps. The dependent record is well-typed by construction in Lean's dependent type theory, and the kernel-only axiom audit confirms no theory-specific axioms are introduced.

The reality certificate as a categorical universal object

The Lean module `Foundation.RealityTerminalCategory` proves a categorical strengthening: every distinguished carrier maps uniquely to its forced reality certificate. In the appropriate category of distinguished carriers, the reality certificate is a terminal object, with a unique morphism from every object.

The structural meaning is that the reality certificate is canonical: any distinguished carrier produces, via the chain, the same reality certificate up to canonical isomorphism. The chain is not merely a possible derivation; it is the universal one, in the technical sense of category theory.

What the master theorem does and does not assert

- + The theorem asserts the existence of an algebraic structure (the reality certificate) that bundles the algebraic forms of spacetime, time, and the constants in recognition-native units, forced from the absolute floor.
- + The theorem asserts that this bundling is canonical (terminal in the appropriate category), so the algebra is forced.
- + The forced algebraic structures are instantiated in observation: $D = 3$ spatial dimensions, a Lorentzian metric of signature $(1, 3)$, a unique massless spin-1 mode (the photon), the speed of causal propagation c , and the constants \hbar and G in the predicted φ -power form (after unit conversion). Each of these is a successful prediction, not a model commitment.
- The theorem does not eliminate the standing demand to check predictions against measurement at higher precision. The numerical band predictions in Section 13 (α^{-1} , η_B) are the current state of that check; the framework’s φ -ladder mass formula and other quantitative predictions are companion work.

The headline

The full chain from $\exists x, y, x \neq y$ to c, \hbar, G is a single theorem in a formally verified library. The premise is the smallest non-trivial propositional commitment; the conclusion is a specific algebraic architecture that observation is found to instantiate.

15 Audit, Honest Scope, Falsification, Open Frontiers

The honest audit

We classify the paper’s claims into four buckets: theorem (proved in Lean), empirical confirmation (a theorem-grade prediction is matched by observation), hypothesis (a numerical band where measurement is asked to land), and model (a genuine interpretive commitment not forced by the chain). The fourth bucket is small; we identify its contents precisely.

Theorem-grade. The forcing chain from the absolute floor through the Aristotelian conditions, the symmetry theorem, the RCL with branch selection, J -uniqueness via Aczél, the Lawvere NNO, time as the canonical orbit, the recognition lattice, $D = 3$ from two convergent forcing arguments, φ from self-similarity, the Lorentzian $(1, 3)$ structure with light cone and energy-momentum relation, the existence of a unique zero-cost massless spin-1 mode with two transverse polarizations, and the constants c, \hbar, G in recognition-native units. The integer -44 in the η_B rung as a structural unification of three combinatorial witnesses. The α^{-1} closed-form expression $44\pi \cdot \exp(-w_8 \ln \varphi / (44\pi))$. The η_B prefactor $c_{RS} = (1 - \varphi^{-8})^2$ at the band level.

Empirical confirmations of theorem-grade predictions

Most of the chain's contact with reality is here. The framework derives a specific algebraic structure as theorem, and observation instantiates it. These are successful parameter-free predictions, not interpretive choices:

1. The framework forces $D = 3$ spatial dimensions. Observation: $D = 3$.
2. The framework forces a Lorentzian metric of signature $(1, 3)$. Observation: the spacetime metric we measure is Lorentzian $(1, 3)$.
3. The framework forces a light cone as the null-displacement locus. Observation: light propagates on a cone with the algebraic form forced.
4. The framework forces a unique zero-cost propagating mode that is massless, spin 1, two polarization states, lightlike. Observation: the photon has exactly these properties.
5. The framework forces the speed of causal propagation $c = 1$ in recognition-native units. Observation: there is exactly one finite causal speed and it is constant in every inertial frame; in SI units it is $c = 299,792,458$ m/s.
6. The framework forces \hbar as φ^{-5} in recognition-native units. Observation: a finite quantum of action exists and matches the predicted φ -power after unit conversion.
7. The framework forces G as φ^5 in recognition-native units. Observation: gravitation has the predicted form after unit conversion.
8. The framework forces $\alpha^{-1} \in (137.030, 137.039)$. Observation: CODATA 2018 reports $\alpha^{-1} = 137.035999084(21)$, inside the predicted band [6].
9. The framework forces $\eta_B \in (6.0, 6.2) \times 10^{-10}$. Observation: Planck 2018 reports $\eta_B = (6.10 \pm 0.04) \times 10^{-10}$, inside the predicted band [16].

These nine items are not interpretive choices. The algebraic chain forces specific structures; observation is found to match. A reviewer who treats item 1 ($D = 3$) as a model commitment is treating Newton's prediction of inverse-square gravity as a model commitment too: in both cases the framework forces a specific structure and observation matches. That is the standard pattern of a successful predictive theory.

Hypothesis-grade

Numerical bands where future precision measurement is asked to land:

- $\alpha^{-1} \in (137.030, 137.039)$ at $\geq 5\sigma$ would be falsified by a measurement outside this band.
- $\eta_B \in (6.0, 6.2) \times 10^{-10}$ at $\geq 3\sigma$ for the prefactor-level prediction; $\eta_B \in (5.5, 7.5) \times 10^{-10}$ at $\geq 5\sigma$ for the rung-level prediction.
- Inter-generation lepton mass ratios on the φ -ladder, deviating at $\geq 5\sigma$ would be falsified.

The framework is parameter-free at the algebraic level, so every numerical prediction is a falsification opportunity.

Genuine model commitments

This is the small bucket the paper is honest about. After audit, we identify exactly two interpretive commitments not forced by the chain, plus the unit-conversion machinery. We name them precisely so a reviewer who wants to challenge the framework on its non-derived content has explicit targets.

1. **The naming and physical interpretation of J .** The cost function $J(x) = (x+x^{-1})/2-1$ is forced by Theorem 6.1 as the unique reciprocal-symmetric, RCL-satisfying, calibrated, continuous solution. We call it *the recognition cost* and treat it as a measure of how distinguishable two ratios are. Mathematically J is also a Bregman-style divergence, an information-theoretic quantity; the framework’s choice to read it as a recognition cost is interpretive. The choice does not affect any numerical prediction (the algebra is the algebra), but it does determine the framework’s vocabulary. A reader who prefers the information-theoretic reading is free to substitute terminology; the chain is unchanged.
2. **The interpretation of recognition-lattice neighborhood as physical adjacency.** The recognition lattice (Section 9) is the quotient of the configuration space by a recognizer’s indistinguishability kernel. The lattice carries a neighborhood relation induced by the recognizer’s finite resolution. We identify this algebraic adjacency relation with physical spatial adjacency. The identification is supported by the dimension and signature theorems (the lattice has the $D = 3$ Lorentzian $(1, 3)$ structure observation finds), but the assertion that algebraic adjacency *is* physical adjacency is a small interpretive step beyond the theorems.
3. **Unit conversions between recognition-native and SI units.** The constants c, \hbar, G in recognition-native units take the values $c = 1, \hbar_R = \varphi^{-5}, G_R = \varphi^5$. To compare with empirical SI values, one needs a unit conversion: how many seconds in a recognition tick, how many meters in a recognition length. These conversions are not model commitments about reality; they are dimensional bookkeeping. We list them as a separate item only because a reader unfamiliar with natural-units arguments may want the bookkeeping made explicit. The bookkeeping does not introduce free parameters: once the recognition tick is matched to one observed period (say, the eight-tick natural period at the chosen energy scale), every other dimensional ratio is fixed.

That is the complete model-content of the chain. Two interpretive choices, both small, neither affecting any numerical prediction; plus one item of dimensional bookkeeping. Earlier formulations of this paper enumerated ten “model bridges” to be conservative; that enumeration was over-inclusive, and we correct it here. Most of what was previously labeled MODEL is more honestly EMPIRICAL CONFIRMATION: the framework derived a specific structure, observation matches, and that is a successful prediction.

Open frontiers

The framework’s frontier register lists open items as of the present writing. Highlights:

Internal gaps. A first-principles Boltzmann-equation derivation of the squared form $(1 - \varphi^{-8})^2$ for the η_B prefactor (the band-level theorem is in hand; the structural argument is interpretive). The exact rung for Ω_Λ . The complete dark-sector mass ladder.

External mathematics axioms. The chain depends on a small number of standard external mathematics facts, used as axioms to avoid recompilation overhead: Bernstein’s inequality (analysis), Hadamard’s classical theorem on entire functions (analysis), Chen’s theorem (number

theory), the S^1 cohomology axiom (used in linking, Section 10). These are uncontroversial. Appendix B provides the audit.

Empirical applications. Applied work in fusion-control, language-model training, and quantum coherence is in preparation as standalone companion papers; these are not detailed here, are not required for any claim of the present paper, and will appear in their own venues.

Falsification strategy

The framework is parameter-free. It can be falsified, sharply, by any of:

- α^{-1} measured outside $(137.030, 137.039)$ at $\geq 5\sigma$;
- η_B measured outside $(5.5 \times 10^{-10}, 7.5 \times 10^{-10})$ at $\geq 5\sigma$ (rung-level falsifier), or outside $(6.0 \times 10^{-10}, 6.2 \times 10^{-10})$ at $\geq 3\sigma$ (prefactor-level falsifier);
- the inter-generation lepton mass ratios $(m_\mu/m_e, m_\tau/m_\mu)$ deviating from the framework's φ -power predictions at $\geq 5\sigma$.

The structural strength of the framework is precisely its falsifiability: with no free parameters, every numerical prediction is a falsification opportunity. The appropriate response to a parameter-free framework that hits its predictions is not skepticism of the predictions but inquiry into how the framework forces them.

What this paper has accomplished

We have walked the forcing chain from $\exists x, y \in K, x \neq y$ to the algebraic architecture of physical reality, in a single self-contained presentation, with epistemic tags on every step and machine-verified backing on every theorem. We have engaged with prior foundational programs and identified specifically what is novel here. We have audited the chain honestly: the bulk of the chain's contact with reality is a sequence of successful parameter-free predictions, with two genuine interpretive commitments and one piece of dimensional bookkeeping the only model content. Two worked numerical examples, plus one mass-ladder illustration, show that the framework's predictions land inside experimental windows. The companion papers cover the technical depth of individual links; the present paper supplies the integration narrative.

A Glossary of Definitions

The glossary is alphabetical. References are to sections of the present paper.

Absolute floor. The single propositional commitment $\exists K, \exists x, y \in K, x \neq y$. Section 2.

Aristotelian conditions. The four classical laws of logic restated as conditions on a comparison operator: Identity, Non-Contradiction, Totality, Composition Consistency. Section 3.

Bit-flip count. On the cube Q_3 , the number of times each of the three coordinate bits flips during the canonical Gray-code recognition cycle, with values $(4, 2, 2)$. Section 13.

Carrier. A type or set K on which the comparison operator is defined.

Comparison operator. A function $C : K \times K \rightarrow \text{Cost}$ assigning a cost to the act of comparing two elements of K . Section 3.

Composition Consistency, (L4). The condition that the comparison operator interacts with composition in the carrier through a single combiner. Section 3.

Cost function J . The unique reciprocal-symmetric, RCL-satisfying, calibrated, continuous cost on $\mathbb{R}_{>0}$: $J(x) = (x + x^{-1})/2 - 1$. Section 6.

CW filtration torsion. The torsion spectrum $\{0, 11, 17\}$ from the cellular filtration of the recognition lattice. Section 13.

Eight-tick. The minimal natural period $2^D = 8$ ticks of the recognition operator at $D = 3$. Section 10.

Eight-tick weight w_8 . The structural sum over per-tick costs of the canonical eight-tick Gray-code cycle. Section 13.

Fermionic DOF. The Standard Model fermionic degrees of freedom, equal to $2 \times D^2(D+2) = 90$ at $D = 3$ via the matter-antimatter doubling. Section 13.

Forcing chain. The dependency-ordered sequence of theorems from the absolute floor to the constants. Sections 2–13.

Golden ratio φ . $(1 + \sqrt{5})/2$; the unique positive solution of $x^2 = x + 1$. Section 11.

Integration gap. The integer $D^2(D + 2)$ governing the rung structure of recognition modes; equal to 45 at $D = 3$. (Called the “consciousness gap” in companion work on consciousness modeling, but used here as a purely combinatorial integer.) Section 10.

Lawvere NNO. The categorical natural-number object characterized by the universal property of having a unique factoring into every other iteration object. Section 7.

Logic-natural-number object LogicNat. The forced arithmetic object on the absolute floor; canonically isomorphic to the Lawvere NNO. Section 7.

Recognition Composition Law (RCL). The functional equation $F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y)$. Section 5.

Recognition lattice. The quotient of the iterated carrier by a recognizer’s indistinguishability kernel. Section 9.

Recognition operator \hat{R} . The iteration step of the framework, with natural orbit Tick. Sections 8, 10.

Recognizer. A surjection $r : \mathcal{C} \rightarrow \mathcal{E}$ from a configuration space onto an event space. Section 9.

Reality certificate RealityCertificate. The bundled output of the master theorem. Section 14.

Recognition-native units. Units in which the forced structural quantities (eight-tick, φ , calibration) take canonical values; constants c, \hbar, G are then φ -powers. Section 13.

Self-similarity. The relation $\sigma^2 = \sigma + 1$ at the carrier identity, forcing $\sigma = \varphi$. Section 11.

Tick Tick. The canonical orbit of the recognition operator, isomorphic to LogicNat. Section 8.

B Machine-Checked Verification and Reproducibility

Repository

The complete formalization is at <https://github.com/jonwashburn/shape-of-logic>. The build command is:

```
elan default leanprover/lean4:v4.27.0-rc1
lake exe cache get
lake build IndisputableMonolith.Foundation.RealityFromDistinction
```

The master axiom audit

The standard query `#print axioms reality_from_one_distinction` returns:

```
'reality_from_one_distinction' depends on axioms:  
[propext, Classical.choice, Quot.sound]
```

These are the standard Lean kernel axioms. There is no theory-specific axiom in the dependency tree of the master theorem.

Reproducibility receipt

For unambiguous verification by reviewers, the following reproducibility data is recorded for the build that backs this paper:

Repository: <https://github.com/jonwashburn/shape-of-logic>

Commit hash: (recorded in repository tag at the time of paper submission; reviewers should consult the latest tagged release named after the paper's arXiv ID)

Lean toolchain: `leanprover/lean4:v4.27.0-rc1`

Mathlib version: (pinned in `lake-manifest.json` of the tagged release)

Sorry count in chain to spacetime and constants: `zero`

Theory-specific axiom count in chain: `zero`

Master theorem axiom dependencies: `{propext, Classical.choice, Quot.sound}`

Build platform: reference build on macOS arm64 and on linux x86_64

Build command: `lake build IndisputableMonolith.Foundation.RealityFromDistinction`

A clean clone of the repository at the tagged commit, followed by the build command above, will reproduce the master theorem and its axiom audit byte-for-byte.

Module map

Selected Lean modules cited in this paper (all in `shape-of-logic`):

Foundation.AbsoluteFloorClosure Joint closure certificate for the absolute floor (Section 2).

Foundation.SelfBootstrapDistinguishability Meta-language self-distinction (Section 2).

Foundation.DistinguishabilityFromSpecifiability Specifiability route (Section 2).

Foundation.NonTrivialityFromDistinguishability Discharge of (L6) from the floor (Section 2).

Foundation.PrimitiveDistinction Aristotelian (L1)–(L3) as type-theoretic facts (Section 3).

Foundation.MagnitudeOfMismatch Symmetry forced by single-valuedness (Section 4).

Foundation.MultiplicativeRecognizerL4 (L4) derivable on multiplicative event space (Section 3).

Foundation.BranchSelection Branch selection: additive branch excluded (Section 5).

Foundation.AlphaCoordinateFixation Coordinate fixation $\alpha = 1$ (Section 5).

Cost.FunctionalEquation.washburn_uniqueness_aczel J -uniqueness via Aczél (Section 6).

Foundation.UniversalForcing.NaturalNumberObject Lawvere NNO and Universal Forcing meta-theorem (Section 7).

Foundation.TimeAsOrbit Time as canonical orbit (Section 8).

Foundation.PreTemporalForcingOrder Pre-temporal forcing order (Section 8).

Foundation.RecognitionLatticeFromRecognizer Recognition lattice from recognizer (Section 9).

Foundation.DimensionForcing Dimension forcing $D = 3$ from linking (Section 10).

Foundation.IntegrationGap Integration gap $D^2(D + 2) = 45$ at $D = 3$, coprimality argument for D odd (Section 10).

Foundation.PhiForcingDerived φ forcing (Section 11).

Unification.SpacetimeEmergence Lorentzian $(1, 3)$ signature (Section 12).

Foundation.ConstantDerivations c, \hbar, G as φ -powers (Section 13).

Cosmology.EtaBExactRungDerivation η_B rung integer -44 via three combinatorial witnesses (Section 13).

Cosmology.EtaBPrefactorDerivation η_B prefactor band (Section 13).

Foundation.RealityFromDistinction.reality_from_one_distinction The master theorem (Section 14).

Foundation.RealityTerminalCategory Categorical universal property of the reality certificate (Section 14).

External mathematics axioms

The Lean library uses a small number of standard external mathematics facts as axioms:

- Bernstein’s inequality (analysis);
- Hadamard’s classical theorem on entire functions of finite order (analysis);
- Chen’s theorem on Goldbach-like representations (number theory);
- The S^1 cohomology axiom used in the linking argument for $D = 3$ (topology).

These are standard external facts; none is theory-specific.

Build status

As of the writing date, the Lean library compiles cleanly with:

- zero `sorry` declarations in the chain to spacetime and the constants;
- zero theory-specific axioms;
- all numerical bounds in this paper proved with closed-form bounds, using Fibonacci identities for φ -power computations.

References

- [1] J. Aczél. *Lectures on Functional Equations and Their Applications*. Academic Press, New York, 1966.
- [2] S. Amari and H. Nagaoka. *Methods of Information Geometry*. Translations of Mathematical Monographs, vol. 191. American Mathematical Society, 2000.
- [3] J. C. Baez and M. Stay. Physics, topology, logic, and computation: A Rosetta Stone. In B. Coecke (ed.), *New Structures for Physics*, Lecture Notes in Physics 813, Springer, 2010, 95–172.
- [4] S. O. Bilson-Thompson, F. Markopoulou, and L. Smolin. Quantum gravity and the standard model. *Classical and Quantum Gravity* 24 (2007), 3975–3994.
- [5] L. M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics* 7 (1967), 200–217.
- [6] E. Tiesinga, P. J. Mohr, D. B. Newell, and B. N. Taylor. CODATA Recommended Values of the Fundamental Physical Constants: 2018. *Reviews of Modern Physics* 93 (2021), 025010.
- [7] A. Connes. *Noncommutative Geometry*. Academic Press, San Diego, 1994.
- [8] L. Crane. Clock and category: is quantum gravity algebraic? *Journal of Mathematical Physics* 36 (1995), 6180–6193.
- [9] B. S. DeWitt. Quantum theory of gravity. I. The canonical theory. *Physical Review* 160 (1967), 1113–1148.
- [10] F. W. Lawvere. An elementary theory of the category of sets. *Proceedings of the National Academy of Sciences* 52 (1964), 1506–1511.
- [11] S. Mac Lane and I. Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer, 1994.
- [12] F. Markopoulou. Quantum causal histories. *Classical and Quantum Gravity* 17 (2000), 2059–2072.
- [13] D. N. Page and W. K. Wootters. Evolution without evolution: Dynamics described by stationary observables. *Physical Review D* 27 (1983), 2885–2892.
- [14] R. Penrose. *The Emperor’s New Mind: Concerning Computers, Minds, and the Laws of Physics*. Oxford University Press, 1989.
- [15] R. Penrose and S. Hameroff. Consciousness in the universe: A review of the “Orch OR” theory. *Physics of Life Reviews* 11 (2014), 39–78.
- [16] Planck Collaboration: N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics* 641 (2020), A6.
- [17] C. Rovelli. Relational quantum mechanics. *International Journal of Theoretical Physics* 35 (1996), 1637–1678.
- [18] C. Rovelli. *Quantum Gravity*. Cambridge University Press, 2004.
- [19] J. E. Shore and R. W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory* 26 (1980), 26–37.

- [20] L. Smolin. *Three Roads to Quantum Gravity*. Basic Books, 2001.
- [21] L. Smolin. *Time Reborn: From the Crisis in Physics to the Future of the Universe*. Houghton Mifflin Harcourt, 2013.
- [22] R. D. Sorkin. Spacetime and causal sets. In J. C. D’Olivo et al. (eds.), *Relativity and Gravitation: Classical and Quantum*. World Scientific, 1991, 150–173.
- [23] M. Tegmark. The mathematical universe. *Foundations of Physics* 38 (2008), 101–150.
- [24] G. ’t Hooft. *The Cellular Automaton Interpretation of Quantum Mechanics*. Fundamental Theories of Physics, vol. 185, Springer, 2016.
- [25] E. Verlinde. On the origin of gravity and the laws of Newton. *Journal of High Energy Physics* 2011 (2011), 029.
- [26] J. A. Wheeler. Information, physics, quantum: The search for links. In W. H. Zurek (ed.), *Complexity, Entropy, and the Physics of Information*. Addison-Wesley, 1990, 3–28.

Recognition Physics Institute preprints (companion papers).

References

- [100] J. Washburn and Z. Zlatanović. Uniqueness of the canonical reciprocal cost. *Axioms* 15 (2026), 1–18.
- [101] J. Washburn. The Law of Logic as a functional equation. Recognition Physics Institute preprint, April 2026.
- [102] J. Washburn. The inevitability of symmetry in logical comparison: magnitude of mismatch as the unique single-valued encoding. Recognition Physics Institute preprint, April 2026.
- [103] J. Washburn. Branch selection in the Recognition Composition Law. Recognition Physics Institute preprint, April 2026.
- [104] J. Washburn. The unified forcing chain: from recognition geometry to the Law of Logic. Recognition Physics Institute preprint, April 2026.
- [105] J. Washburn. The Universal Forcing meta-theorem: one Law of Logic, one arithmetic, across all realizations. Recognition Physics Institute preprint, April 2026.
- [106] J. Washburn. The cosmological forcing chain: from distinction to the arrow of time. Recognition Physics Institute preprint, April 2026.
- [107] J. Washburn. Observer forcing in Recognition Science. Recognition Physics Institute preprint, April 2026.
- [108] J. Washburn. Geometry from logic: a comprehensive treatment. Recognition Physics Institute preprint, April 2026.
- [109] J. Washburn. Closure of the absolute floor for the Recognition Science forcing chain. Recognition Physics Institute preprint, April 2026.
- [110] J. Washburn. Self-bootstrap distinguishability: the meta-language already distinguishes. Recognition Physics Institute preprint, April 2026.

- [111] J. Washburn. Distinguishability from specifiability. Recognition Physics Institute preprint, April 2026.
- [112] J. Washburn. The φ -ladder mass formula and the Standard Model fermion spectrum. Recognition Physics Institute preprint, April 2026.
- [113] J. Washburn. Shape of Logic: Lean 4 formalization. Public repository, 2026. <https://github.com/jonwashburn/shape-of-logic>.