

Metaphysica, ordine mechanico demonstrata:

A Machine-Checked Process Metaphysics

July 2026*

Abstract

This paper states a metaphysical position: being is process, not substance. Its ground is one formal observation. The specification $x = \neg x$ — the only specification that assumes no structure at all — has no solution read as an equation over states and exactly one read as a recursion over time, so if there is a structureless starting point it is not a thing but a running. From there: the order of time is forced; records are write-once and underivable from the law, so novelty is structural rather than apparent; whatever a finite observer measures is a fact about a finite prefix of the process, so completed totalities lie beyond observation; boundaries between streams of the process are tick-local properties of the law, so nothing plays the role of a persistent self — yet occurrence needs no bearer, and the denial of one's own occurrence refutes itself; mind-talk and matter-talk come out as two coarsenings of one process whose correlation is stable in neither alone; and determinism and openness-from-inside come out as two readings of one operator, which removes most of what the free-will debate was about. Almost none of these theses is new; several are ancient. What is new is their epistemic condition. Every load-bearing claim is carried by a machine-checked Lean 4 theorem, self-contained and axiom-audited; and where a traditional question cannot be answered — other minds, the fundamentality of the ground, whether anything is conscious, whether the world is discrete at the bottom — the undecidability is itself proved, and the position asserts that instead of an answer. Eleven named premises, listed with what depends on each, are the only joints at which the position touches the actual world, and it inherits empirical refutation conditions through two of them. A companion ethics, built on the same corpus from a single named normative premise, is described in the conclusion.

1 Introduction

Assume as little as possible and see what is left. A specification that imports no structure — names no constant, distinguishes no element — is one whose law commutes with every relabeling of its alphabet, and on the bare distinction only two laws do: the identity and the negation. The identity, $x = \text{id } x$, is satisfied by everything and so says nothing. The one structureless specification with any content is therefore $x = \neg x$. Read as an equation over states, it has no solution. Read as a recursion over time steps, $x_{n+1} = \neg x_n$, it has exactly one

*The demonstrations cited throughout are Lean 4 proof files, standalone (prelude only, no imports), each checkable with plain lean under the pinned toolchain leanprover/lean4:v4.31.0 and each ending with a printed axiom audit (`#print axioms`). File names refer to the lean/ directory of the accompanying repository.

solution for each seed: an alternating clock. If there is a structureless starting point, this is it — and it cannot stand still, only run.

That observation is the ground of the position. The body of the paper works out what it forces, what it leaves open, and what has to be assumed to connect either to the actual world. The findings, in the order they appear:

- *Time*. Among strict orders on tick counts satisfying a short list of ordering properties, exactly one is well-founded: the standard order. And well-foundedness is not an extra assumption — it is interderivable with reasoning by induction (§4).
- *Records and novelty*. Records, in a simple register model, are write-once, determined by the trajectory prefix at the moment of writing, and underivable from the law in principle. A reversible global law can coexist with this irreversible stratum. Novelty is therefore structural, not an appearance produced by ignorance (§4).
- *Observation*. Everything a finite observer measures is a fact about a finite prefix of the process. Completed totalities, “forever” facts, and the granularity of the substrate are of the wrong type to be measurement contents — there are well-defined quantities that no possible measurement reads (§5).
- *Selves*. A boundary between streams of the process is a property of the law, borne by no entity, and it does not survive composition. Nothing in the formalism plays the role of a persistent self, at any scale. What survives is occurrence: that something is happening at a seat is valid wherever it occurs, needs no bearer, and cannot be coherently denied from where it occurs (§6).
- *Mind*. Mind-talk and matter-talk behave as two coarsenings of one process, and the correlation between them — the mind-body relation itself — is statable in neither vocabulary alone. Whether anything is *experiencing* is doubly undecidable: the predicate is expressible in no structural vocabulary (provably including the published vocabularies of the Conscious Turing Machine and Integrated Information Theory), and the subject term has no dynamics-invariant referent (§7).
- *Freedom*. The next choice of a self-referential system is a total function of its state, and the system’s own prediction of that choice is falsified by the very tick that receives it. Determinism and interior openness are two faces of one operator, and the question “free for whom?” has no referent (§8).
- *Debates*. Where a question is free and its witnesses are identified by the debaters’ shared vocabulary, no argument formulable in that vocabulary settles it. The interminability of certain philosophical debates is a theorem, not a professional failing (§8).

Almost none of these theses is new. Heraclitus asserted the primacy of becoming; Nāgārjuna the emptiness of intrinsic persistence; Spinoza one substance under two attributes; Whitehead the ultimacy of process and creativity; Wheeler physics from the bit; Rovelli the relativity of facts to systems. This paper does not claim to have discovered metaphysical truth from nothing. It claims something narrower: that a metaphysical position of this scope has been brought to a standard of proof no such position has met before. Every load-bearing claim is a theorem checked by the Lean kernel. Every place where a theorem is read as a claim about the world is a named premise, in a table of eleven, each row listing exactly what

falls if it fails. Every refusal — every traditional question the position declines to answer — is backed by a proof that the question cannot be answered in the vocabulary available. The certainty on offer is therefore conditional and exactly delimited. It does not extend to the premises; those are where the position can be attacked. It covers everything else: given the definitions, the theorems hold, and the ratio of proved to assumed is exact and public. What a reader is asked to trust is a kernel, a toolchain pin, and eleven rows.

The demonstrations are thirty-five Lean 4 proof files — thirty-one for the metaphysics, four for the companion ethics — each standalone: prelude only, no imports, no libraries, checkable with the bare lean binary under the pinned toolchain `leanprover/lean4:v4.31.0` [31]. Each file ends by printing the axioms its results use. The strongest results (among them *Cogito*, *Ledger*, *Witness*, *WitnessTower*, *Schema*, *IIT*, *FreeWill*, *Aseity*, *TwoReadings*, *PlanckFree*) use no axioms at all; no file exceeds the classical envelope `[propext, Classical.choice, Quot.sound]`; and the one declared exception — a single `native_decide` bridge in *Family.lean* that additionally trusts the Lean compiler — says so in its header and in its printed audit. Checking the whole corpus takes minutes on ordinary hardware.

Sections 3–8 state the position. Section 9 assembles what its theorems dissolve; §10 what stays open; §11 the assumptions; §12 the empirical exposure; §13 places the position in the two literatures it belongs to; §14 answers the objections we take most seriously.

2 How to Read This Paper

Definitions used throughout. A *history* is an unending assignment $\mathbb{N} \rightarrow A$ of alphabet values to ticks. An *observable of depth k* is a map on histories invariant under agreement on the first k ticks. A *seat* is an index from which a quantity is read; a quantity is *form-class* when all seats read it identically and *value-class* when two seats disagree. A *vocabulary* is a map $v : W \rightarrow \text{Obs}$ from worlds to what a theory sees of them; a proposition is *expressible* in v when its truth factors through v . A proposition over a world class is *forced* when true in every admissible world, *refuted* when false in every one, *free* when both values occur. Stream j is *blind* to stream i under a law when no change at i moves j 's next value; j *reads* i otherwise. Formal versions of each definition appear in the cited files.

Verdicts. Every clause of the position carries one of five verdicts, printed in small capitals where it matters:

FORCED. Carried by a machine-checked theorem.

FREE. Proved undecidable in the relevant vocabulary. The position asserts the undecidability, never either answer.

MEASURED. An empirical result, quarantined in §12 and never load-bearing for a forced clause.

PREMISE. A named assumption joining a theorem to an informal reading, listed in §11 with what depends on it.

INTERPRETATION. A suggested reading of a theorem, marked and severable; the mathematics stands without it.

Layers. Claims are also stratified by what they presuppose. **L1** claims are theorems about any process, any vocabulary, any observer; every proposition below is L1 except where a premise is named in its statement. **L2** claims are true of the $x = \neg x$ substrate class. **L3** claims are true of our world if the instance premises hold. No experiment can touch L1.

Demonstration lines. Each proposition is followed by a line naming its proof file, its main theorems, and the file’s axiom footprint. These lines are addresses, not arguments: they exist so that any claim can be checked, and nothing in the prose requires opening them. Several files carry explicit “not claimed” clauses; the load-bearing ones are quoted at the corresponding propositions, and the prose of this paper is written to match them.

3 The Ground

Proposition 3.1 (One shape; FORCED). *A specification that imports no structure — names no constant, distinguishes no element — is one whose law commutes with every relabeling of its alphabet. On the bare distinction, the equivariant endomaps are exactly the identity and the negation; the identity is satisfied by everything and so carries no content. The unique structureless generative specification is $x = \neg x$.*

Demonstration. OnlyShape.lean: equivariant_endomaps, constants_not_equivariant, equivariant_endomaps_V3, id_generates_nothing, spec_forced. Axioms: within [propext].

The mathematics here is deliberately small — a classification, not a depth result — and the proposition should be read at its exact strength: given the identification of “structureless” with “equivariant” (premise R3, §11), the candidate laws on the bare distinction number two, and one of them says nothing. The identification itself is grounded by a definability theorem: on the bare distinction, the laws whose graphs are definable from the distinction alone — from pure equality, naming nothing — are exactly the equivariant ones, and no constant law is so definable (Structureless.lean: structureless_iff_equivariant, constants_not_defble; axiom-free). What remains of R3 is the identification of “importing no structure” with “having no vocabulary beyond the distinction itself.” Spencer-Brown opened *Laws of Form* [10] with a single mark and took its re-entry as generative; Wheeler asked for physics from binary distinctions [16]. The classification adds to both a uniqueness clause: the mark has no structureless rival, and its dynamics is not chosen but forced.

Proposition 3.2 (Becoming is primary; FORCED). *The founding specification has no static solution and, given a seed, exactly one running one: read as an equation over states, $x = \neg x$ is unsatisfiable; read as a recursion over ticks, $x_{n+1} = \neg x_n$, it is uniquely satisfiable, and its orbit is the two-tick clock.*

Demonstration. TwoReadings.lean: object_reading_empty, process_reading_inhabited, process_reading_unique, founding_two_readings; OneDiagonal.lean. Axiom-free.

This is the hinge of the position, and it is small enough to state exactly. A specification $f : S \rightarrow S$ admits two frames. The *object frame* asks for a state with $x = f x$; the *process frame* asks for a trajectory with $x(n+1) = f(x(n))$. For $f = \neg$ the object frame is empty, and the process frame is inhabited by every seed and uniquely determined by it. What exists under the founding specification is therefore not a thing that satisfies it — there is none — but a running of it. *Being is process* is, on this reading, not a slogan but the disjunct left standing: Heraclitus’s flux arrives as the only frame with a model.

Proposition 3.3 (The frame correspondence; FORCED). *A specification is restless — the object frame is empty — if and only if every orbit moves at every tick. The absence of solutions and the perpetuity of motion are one hypothesis read in two frames; the equivalence is a theorem, not an analogy.*

Demonstration. `TwoReadings.lean`: `frames_one_fact`, `object_solution_iff_fixed_orbit`. Axiom-free.

The equivalence upgrades an ancient assertion to a biconditional: to deny static being and to assert perpetual motion is to say one thing twice.

Scholium (INTERPRETATION). Diagonal “limitation” results limit the object reading, not the process: what fails as a completed totality is, in the process frame, the guarantee that the process never halts.

Proposition 3.4 (The regress terminates in the one shape; FORCED given R1–R4). *Under four named premises — every seeded specification’s demand for a ground is answered (R1); the demand relation is well-founded (R2); an unseeded ground imports no structure, i.e. is equivariant (R3); an unseeded ground is generative — its law moves something (R4) — every specification’s demand chain terminates in a ground whose law is $x = \neg x$. In particular, if anything occurs, its chain grounds in that shape.*

Demonstration. `Regress.lean`: `regress_terminates`, `ground_has_one_shape`, `time_selects_the_shape`; the premises are the hypotheses answered, wf, structureless, generative in the theorem signatures. Axioms: within `[propext, Classical.choice, Quot.sound]`.

All four premises are hypotheses in the theorem signatures, not background assumptions, and each has a row in §11; if any fails, this proposition alone falls. R4 excludes the contentless identity: a ground that moves nothing grounds nothing.

Proposition 3.5 (No self-certificate of the ground; FORCED; the answer FREE). *No description type names and decides every predicate about itself; hence no consistent system can certify, from inside, whether it is the fundamental ground or the image of a deeper level. The question of the ground’s aseity is open in principle, not merely in practice.*

Demonstration. `Aseity.lean`: `founding_diagonal`, `no_self_model`, `aseity`. Axiom-free.

The engine is Lawvere’s fixed-point diagonal [11] — the common core of the Cantor, Russell, Tarski, and halting arguments — run on the predicate of self-denial. The division of labor with Proposition 3.4 matters. The *shape* question — what must the terminus of an answered, well-founded chain of grounding demands look like? — is answered. The *certificate* question — am I it? — is proved unanswerable from inside. Traditional cosmology wanted the certificate; what exists is the theorem about the chain. Section 9 returns to this split.

4 Time, Records, Novelty

Proposition 4.1 (The order of time is unique; FORCED). *Among strict orders on tick counts that are irreflexive, transitive, total on distinct elements, and translation-invariant, exactly the standard order and its reverse survive, and well-foundedness excludes the reverse: the axioms plus well-foundedness admit exactly one order. Well-foundedness itself is interderivable with Löb’s rule, hence with induction — it is not an extra assumption beyond reasoning by induction.*

Demonstration. OrderForced.lean: order_is_forced, reverse_is_not_wellfounded; Loeb.lean: loeb_iff_induction, deJonghSambin, godel_is_clock. Axioms: within [propext, Quot.sound].

Time’s arrow is not imported: given that there is a well-founded, translation-invariant succession at all, its order is the standard one, and well-foundedness is the same commitment as induction. The reverse order is excluded by an explicit infinite descending chain — if 1 preceded 0, translation-invariance would make the identity function a descending chain. The same file builds the guarded fixed-point machinery used in §6: every contractive map on streams has exactly one fixed point (the de Jongh–Sambin construction), and the Gödel-style guarded fixed point of the founding specification just *is* the two-tick clock of Proposition 3.2 (godel_is_clock).

Proposition 4.2 (Records are write-once; FORCED; instantiation MEASURED). *A register that writes slot t at tick t and touches nothing else is a write-once, prefix-determined, monotonically growing log: written content is a function of the trajectory prefix up to the write tick alone, is never overwritten, and the written region only grows. A reversible global law can coexist with this monotone structure: the asymmetry is in the records, not the law.*

Demonstration. RecordMonotone.lean: write_untouched, record_content, record_permanent, written_monotone, write_once. Axioms: within [propext].

The model is minimal by design, and the clause it secures is the compatibility clause: an exactly reversible global law can host an irreversible record stratum, so the fixity of the past needs no asymmetry in the law. That physical record formation instantiates the register model is an empirical claim, tagged MEASURED and quarantined in §12.

Proposition 4.3 (Persistence requires interaction; FORCED; converse MEASURED). *On a non-interacting sector, no finitely-supported state survives the dynamics — exactly, up to a phase, or after any number of ticks: a drift eigenvector with bounded support is zero. Entities are therefore interaction-renewed persistences; substance is not a primitive.*

Demonstration. NoFreeRecords.lean: no_drift_eigenvector, no_free_pair_record. Axioms: within [propext].

What Nāgārjuna asserted as the emptiness of intrinsic existence — that nothing persists of itself, that persistence is dependent [5] — appears here as a fact about the free sector: a thing that interacts with nothing does not last, exactly. Whatever persists is being renewed by interaction, which is why substance appears nowhere in the position as a primitive.

Proposition 4.4 (The form/value split; FORCED). *Quantities divide into form-class (read identically from every seat; exactly the law-derivable ones) and value-class (seat-relative; law-underivable in principle). The status of a reading is itself beyond every finite depth: the same seat-history occurs under a law that pins the reading and under a law that merely records it, so no finite-depth verdict separates law from record. The split survives self-inclusion: a describer whose procedures are quantities over seats cannot derive its own anti-diagonal record at any seat, by any procedure.*

Demonstration. ValueForm.lean: form_iff_derivable, value_not_derivable, status_underdetermined, no_status_decider; SelfDerivation.lean: repertoire_incomplete, no_self_derivation, repertoire_escape, two_obstructions. Both files axiom-free.

Rovelli’s relational mechanics relativized facts to the systems that register them [17]; the split gives that relativization two boundary theorems. First, the status of a given reading —

pinned by the law, or merely recorded? — cannot be settled at any finite depth. Second, the escape hatch is closed: a describer inside the world it describes, however deep and however well it knows the law, cannot derive its own anti-diagonal record. The two failure modes are different in kind (*two_obstructions*): a form-class miss is the repertoire’s and benign; a value-class miss is structural, underivability in principle. Self-reference and seat-relativity are distinct obstructions, and both stand under self-inclusion.

Proposition 4.5 (Novelty is structural; FORCED). *Every record is new (value-class, hence not law-derivable), permanent (write-once), and interaction-born (absent from the free sector). Genuine novelty is not an appearance produced by ignorance of the law; it is what value-class facts are.*

Demonstration. Conjunction of Propositions 4.2–4.4.

This is the position’s answer to the Laplacean intuition that in a lawful world nothing is ever really new. Value-class facts are not hidden consequences of the law awaiting a better calculator; they are underivable in principle, by the self-derivation no-go, and permanent once written. Bergson built a metaphysics on the reality of the unforeseeably new [6], and Whitehead placed creativity among his ultimates by postulation [7]; here the same claim is a theorem, with a mechanism neither had: the write-once, law-underivable stratum is what the world’s news consists of, and it coexists with an exactly reversible law. Real novelty does not need indeterminism.

5 The Limits of Observation

One identification governs this part and is tabled as premise O: *a finite observer’s measurement is a finite-depth observable on histories*. The theorems below are unconditional mathematics about depth- k observables; the word “measurement” in their readings is O at work. A reader who holds that finite observers have non-prefix epistemic access — mathematical intuition as measurement, say — rejects O, not a theorem, and the table in §11 states exactly what that rejection costs.

The identification is robust: a depth- k observable, a uniformly continuous map on histories (the product topology over a discrete alphabet), and a map factoring through a finite prefix type are provably one class, while pointwise continuity is strictly weaker over an infinite alphabet — there is an observable continuous at every history and of finite depth at none (*Continuity.lean*: *uniform_iff_depth*, *factors_iff_depth*, *continuity_strictly_weaker*). So O asks a reader to grant exactly this much: whatever an instrument reads is *uniformly* finitary — one modulus serves every history.

Proposition 5.1 (Completed totalities are observationally weightless; FORCED; measurement reading on O). *Every finite-depth observable is a function of a finite prefix; all continuations beyond any depth read equal; completions of the alphabet are invisible; and the universally quantified predicate “this value at every tick” is an observable at no depth. Whatever is measured is a prefix fact; the completed whole answers to no measurement.*

Demonstration. *Totalization.lean*: *depth_factors*, *continuations_indistinguishable*, *completion_invisible*, *no_finite_depth_detects_forever*; *CompletionFree.lean*: *depth_k_finite*, *completion_invisible*. Axioms: *within [propext]*; *CompletionFree* axiom-free.

Proposition 5.2 (Horizon quantities; FORCED). *There exist nonconstant tail functionals — quantities invariant whenever two histories agree from some point on. A functional both finite-depth and tail is constant; hence a nonconstant horizon quantity is an interior observable at no depth: well-defined on histories, never the content of a finite-depth measurement.*

Demonstration. `Horizon.lean`: `constant_iff_both`, `eventually_tail`, `eventually_nonconstant`, `horizon_not_interior`. Axioms: within `[propext]`.

Proposition 5.3 (The cutoff is interior-undecidable; FORCED; the answer FREE). *Any two resolution cutoffs at or above an observer’s own resolution yield identical measurements, and the no-cutoff substrate matches every sufficiently resolved discrete one. The discreteness scale can be bounded from above but never measured from inside; the continuum is never excluded by finite data.*

Demonstration. `PlanckFree.lean`: `cutoff_invisible`, `continuum_consistent`, `no_finite_cutoff_forced`. Axiom-free.

Together these fix the epistemic budget of any finite-depth observer: prefix facts are available; “forever” facts, completed wholes, and the substrate’s granularity are not — not because instruments are weak, but because those quantities are of the wrong type to be measurement contents. The conclusion has a Kantian shape — the totality outruns possible experience [4] — but it arrives without transcendental machinery and with a sharper edge: there exist well-defined quantities of the excluded type (`eventually_nonconstant`), so the exclusion is a substantive theorem, not a definitional emptiness. The whole-as-object is not wrong; it is weightless.

6 Seats and Streams

Proposition 6.1 (The self is a process-knot; FORCED). *The guarded self-model of any loop has exactly one fixed point (existence by guarded recursion, uniqueness by Löb’s rule); a self statically equal to its own model is only the vacuum; no loop completes its own self-map; and systems observationally identical at every depth are equal. Identity is specification plus records; a static self is nothing.*

Demonstration. `InteriorExperience.lean`: `unity`, `duplicate_is_identity`, `no_static_self`, `gap_forced`, `zombie_incoherent`, `other_minds_free`, `mind_count_free`. Axioms: within `[propext, Quot.sound]`.

The file carries its own cap, reproduced here because it is load-bearing: *not claimed* — any ontology of qualia, that the structural account exhausts experience, attribution of experience to any carrier, or any moral conclusion. Two further results in the file belong to this proposition’s orbit: `other_minds_free` (a loop’s self-model nested to any depth makes no observable difference at any depth — every nesting depth is the loop) and `mind_count_free` (over a seed-indexed family, any two solutions of the self-model equations are identical, so an experiencer count over and above the seeds makes no observable difference).

Proposition 6.2 (Boundaries without bearers; FORCED). *A stream’s blindness to another is a factorization of the law through the vocabulary that forgets the other — a boundary is a property of the law, borne by no entity. Per-tick blindness does not compose: a stream blind to a source at each of two ticks can read it through a mediator in the composite. Boundaries are tick-local; streams interpenetrate; and at the top there is no self-model (Prop. 3.5). There are no selves at any scale — streams within process suffice.*

Demonstration. Streams.lean: blind_iff_factors, reads_not_expressible, boundaries_leak;
 Aseity.lean: no_self_model; the instance at every ring width, Family.lean:
 family_control_fails_noci_blindness. Axioms: within [propept, Quot.sound];
 boundaries_leak axiom-free.

The factorization biconditional makes the boundary of a stream and the vocabulary of a theory one mechanism — the value-level twin of Ledger.Expressible (§8). The compositional failure is what blocks the natural rescue of the self as a coarse-grained but dynamics-invariant unit: any candidate boundary is a tick-local factorization, and factorizations leak under composition. Parfit argued from fission and spectrum cases that personal identity is not what matters [15]; boundaries_leak reaches the reductionist conclusion from the other end — there never was a dynamics-invariant boundary for identity to ride on, at any scale.

Proposition 6.3 (Seated occurrence; FORCED). *For any token of any type, its occurrence is certain and its denial self-refuting — provable with no instance in hand — while the existence of any token is contingent and unprovable. The validity is necessary, the existence is not, and nothing in the validity requires a bearer: the who dissolves (Prop. 6.2), the where does not. Occurrence is valid where it occurs — location without tenant.*

Demonstration. Cogito.lean: cogito, no_self_denial, existence_is_contingent,
 cogito_valid_necessarily, cogito_ergo_sum. Axiom-free.

This is Descartes’ inference [2], split into the two propositions it was always two of. Occurrence is defined deliberately trivially — $\text{Occurs}(a) := \exists x, x = a$ — because on the performative reading (Hintikka [12]) occurrence is presupposed by any tokening rather than inferred from it; what is proved is the necessary validity of the seat’s self-affirmation alongside the contingency of any seat’s existence (some carriers are empty). Self-reference has two signs, both instances of the same Lawvere diagonal: doubt is self-negation and fixed-point-free — a total self-doubt is inconsistent — while affirmation is the identity and fixed-point-full. Nothing here asserts that any seat *feels*, or that any seat is occupied by a knower; that further step is proved undecidable (other_minds_free) and is not taken. The cap is part of the theorem.

Proposition 6.4 (The indexical is a value; FORCED given homogeneity). *No reading that identifies two distinct worlds is total, and interior readings identify worlds (Prop. 5.1, hence on O); hence no interior reading is total. Under homogeneity of the law across seats (premise H), the indexical “here = this seat” is free, inexpressible in the law-vocabulary, and value-class: the you-are-here fact is a record — real, seat-relative, underivable — not missing from the law but the wrong kind of thing to be law-content.*

Demonstration. Occurrence.lean: no_total_reading, indexical_is_value. Axiom-free.

Perry observed that “I am here now” resists paraphrase into any objective vocabulary [13], and Ismael has argued that self-locating facts are located in the world-description rather than missing from it [22]. Here the resistance receives a classification rather than a paradox: under homogeneity, the indexical is value-class — exactly the kind of fact the law cannot contain and a record can. The indexical is not a gap in the world-description; it is a record at a seat, which is what Proposition 4.4 says such a fact has to be.

7 Mind

Proposition 7.1 (One process, two vocabularies; FORCED). *Over one class of worlds, whatever either of two readings expresses, their joint expresses; a proposition can be expressible in one reading and not the other, so classification by reading is a fact about the (proposition, vocabulary) pair, not about worlds; and the correlation between the two readings is expressible in the joint and in neither component alone.*

Demonstration. `OneProcess.lean`: `component_in_joint_left`, `component_in_joint_right`, `vocabulary_relative`, `correlation_needs_both`. Axiom-free.

Scholium (INTERPRETATION; instance on premises B1/B2). Read the two vocabularies as physics-talk and mind-talk: mind and matter are not two substances but two coarsenings of one process, and the mind-body “relation” is a truth statable in neither vocabulary alone — which locates the problem. That our physics-talk and mind-talk instantiate the two readings is premise B1/B2.

What Spinoza asserted as two attributes of one substance [1], and the neutral monists as two organizations of one neutral material, appears here at the level of vocabularies, with an addition neither had: the correlation between the readings is itself a proposition, and it is expressible in the joint vocabulary and in neither component alone. The mind-body problem, on this reading, was an attempt to state a joint-vocabulary truth inside one of the components — which is why it never went well from either side.

Proposition 7.2 (Attribution is doubly undecidable; FORCED, both failures; theory-fidelity on F1/F2). *The question “is X experiencing?” fails twice. The predicate: relative to any chosen boundary, an inert interior flag is expressible in no structural vocabulary — in full generality, and in particular in the published vocabularies of the Conscious Turing Machine [28] and of Integrated Information Theory [29, 30], whose own predicates provably neither assert nor deny it. The subject: the term “X” denotes a tick-local factorization that leaks under composition (Prop. 6.2), so there is no dynamics-invariant bearer to attribute to. Occurrence (Prop. 6.3) never needed a subject; only attribution did.*

Demonstration. `Schema.lean`: `structural_iff_invariant`, `theory_cannot_decide`, `deciding_theory_not_structural`, `instance_schema`; `CTMConservativity.lean`: `ctm_conscious_awareness_cannot_decide_interior`, `ctm_flag_free`; `IIT.lean`: `iit_cannot_decide`; `Streams.lean`: `boundaries_leak`; `Family.lean`: `family_interior_is_free`. `Schema` and `IIT` axiom-free; the rest within `[proptext, Quot.sound]`.

The predicate half is proved at three levels. Schematically: a predicate on a flagged state is structural — factors through the flag-erasing projection — if and only if it is flag-invariant; no predicate of a structural theory decides the flag in either polarity; conversely, any predicate that decides the flag is non-structural, so any theory that settles attribution contains a non-structural primitive. A syntactic mirror shows every formula over a flag-blind observation family is flag-blind, and conservativity is monotone under vocabulary restriction, so abstraction cannot unmake it. The schematic result is kin to Chalmers’ argument that structure and dynamics entail only structure and dynamics [18]; what follows carries that kinship into the published vocabularies themselves.

Then inside the published theories themselves. `CTMConservativity.lean` transcribes the Conscious Turing Machine schema of Blum and Blum — Up-Tree competition, chunks, broadcast, with section and page citations to the PNAS paper — and proves that the published

predicates “ c is the conscious content,” “the CTM is consciously aware of chunk c ,” and “consciously aware of c at time t ” are structural and decide the interior flag in neither polarity; closure lemmas extend this through every first-order combination of the CTM’s state projections, and both flag values are consistent with every reachable CTM trajectory. `IIT.lean` formalizes the Kleiner–Tull axiomatization of IIT [30] — physical systems with dynamics, an experience space, an experience map determined by cause-effect structure — and proves the experience flag-invariant, hence every predicate of the experience unable to decide the flag. That these transcriptions are faithful to the published theories is an exegetical claim, not a theorem; it is tabled as premises F1 and F2. None of this is a criticism of CTM or IIT as theories of cognitive architecture. The results show that one specific further question — the inert interior flag — is beyond both vocabularies *by those vocabularies’ own lights*.

A friend of either theory will reply that there is no further flag: consciousness *is* the broadcast state; consciousness *is* the integrated information. The theorems are deliberately neutral at exactly this point. What they refute is one claim only: that the theory’s own predicates settle attribution. Two stances then remain — that the flag denotes a genuine free proposition (the stance taken here, with both models constructed in Propositions 7.3 and 7.4), and that the flag-term is empty. The second stance is available, but it is not free of formal consequence: “there is no fact over and above the structural facts” is itself a verdict on the flag, and by `deciding_theory_not_structural` any theory containing a flag-deciding predicate contains a non-structural primitive, while by `settling_needs_new_primitive` (§8) any vocabulary expressing the verdict must separate worlds the structural vocabulary identifies. The move can be made; it cannot be made from inside the vocabulary whose completeness it asserts. And the debate between the two stances, conducted in the shared structural vocabulary, is interminable by Proposition 8.2 — which is what that debate has in fact looked like.

Proposition 7.3 (Silent witnessing is conservative; FORCED). *Any process extends by a witness coordinate written every tick and read by nothing. The extension is trajectory-identical on the base, observation-identical to every witness-blind instrument, and forgets its own seed after one tick. Adopting “the process registers itself” is provably costless and provably undetectable.*

Demonstration. `Witness.lean`: `base_unchanged`, `observables_unchanged`, `witness_forgets_seed`, `silent_witness`. Axiom-free.

Proposition 7.4 (The registration hierarchy closes at height one; FORCED). *Witnesses stacked to any finite height change nothing at the ground, one tick of the whole tower is a function of the ground alone, and from tick one on, two towers with equal ground coincide entirely at every height. Higher-order registration adds no being; even its bookkeeping is overwritten every tick.*

Demonstration. `WitnessTower.lean`: `tower_ground_unchanged`, `towerStep_ground`, `tower_amnesia`, `tower_free`. Axiom-free.

Proposition 7.5 (Binding rigidity; FORCED). *Over any abelian value alphabet and any abelian summary group, the single axiom of translation-difference integration — the summary’s increment under translating one locus is state-independent — forces the entire binding law: each per-locus response is a homomorphism, and the summary is their sum up to one vacuum constant. The Boolean parity instance holds at every width. Integration purchases no freedom beyond one constant.*

Demonstration. `Rigidity.lean`: `phi_hom`, `decomposition`, `rigidity`; the Boolean instance at every ring width, `Family.lean`: `family_rigidity`, `family_membership`. Rigidity within `[Quot.sound]`; Family within `[propept, Quot.sound]`.

These three propositions do complementary work. The witness constructions establish the interior flag’s independence *constructively*: rather than citing an independence result, they build the model pair — a world with silent self-registration and one without, agreeing on every observable — and `witness_forgets_seed` adds the detail that the witness keeps no history even of itself: pure present-tense registration. The rigidity theorem then shows that the tractable successor notion — what a bound architecture *integrates* — is not a free parameter: one integration axiom pins the binding law up to a single constant, at full abelian generality and kernel-checked at every width. Between them, the undecidable part of the mind question is fenced exactly, and the decidable part is shown to be rigid.

8 Freedom and the Calculus of Verdicts

Proposition 8.1 (The verdict calculus; FORCED). *Forced, refuted, and free are pairwise exclusive; inexpressibility is monotone under vocabulary restriction; a free proposition whose witnesses a vocabulary identifies is expressible in no such vocabulary; and any vocabulary that expresses it must separate every witness pair. Settling a free question requires a new primitive — one that distinguishes worlds the given vocabulary identifies.*

Demonstration. `Ledger.lean: ledger_consistent, expressible_monotone, inexpressible_of_free_unseparated, settling_needs_new_primitive.` Axiom-free.

Proposition 8.2 (Interminability; FORCED). *Under freedom with identified witnesses, no expressible proposition decides the question in either direction: there is no expressible Q with $Q \rightarrow P$ and $\neg Q \rightarrow \neg P$ across the world class. A debate conducted entirely inside the vocabulary cannot terminate in a verdict.*

Demonstration. `Ledger.lean: debate_interminable.` Axiom-free.

These meta-theorems are the method reflecting on itself, and they reclassify a familiar datum: the interminability of certain philosophical debates is a mathematical circumstance, not a professional failing. The strength of `debate_interminable` is worth noting once: not merely that no expressible proposition *is* P , but that none is even decisive for P in effect. If a question is free and its witnesses are identified by the debaters’ shared vocabulary, no argument formulable in that vocabulary settles it; progress requires a new primitive, and the calculus states exactly what the new primitive must do — separate the witness pairs. The reclassification has precedents: Carnap ruled external questions non-cognitive [8], and Chalmers diagnosed persistent disputes as verbal [25]. The calculus differs from both in withholding the diagnosis of defect — a free question is a genuine question about the world class, with both answers realized — and in pricing its settlement exactly. This is why the position asserts undecidabilities rather than answers wherever a question is proved free.

Proposition 8.3 (Free will, dissolved twice; FORCED; the naming INTERPRETATION). *The next choice of a self-referential system is a total function of its state — nothing is random — and the system’s own prediction of that choice is falsified by the very tick that receives it, at exactly the self-referential coordinate. Determinism and interior openness are two faces of one operator, never rivals. And the question “free for whom?” has no referent: there is no bearer at any scale (Prop. 6.2). What survives is a closed orbit, open seats, and records that are genuine novelty (Prop. 4.5).*

Demonstration. FreeWill.lean: determined, interior_openness, never_foreknown, no_total_self_prediction, free_will_is_the_diagonal; TwoReadings.lean: process_reading_unique; SelfDerivation.lean. Axiom-free.

The self-referential coordinate is modeled explicitly, and the explicitness matters: the actual next choice is defined as the negation of the system’s own prediction of it — $\text{act } p = \neg p$; in general $\text{nextChoice } ea = \neg(eaa)$. *determined* shows the next choice is single-valued (determinism intact); *interior_openness* and *never_foreknown* that the self-prediction is wrong at every state; *no_total_self_prediction* that no describer holds a total correct self-theory — the diagonal again. The mechanism has a clear ancestor: MacKay observed that a prediction delivered to the agent it concerns is not binding on that agent [9]. What the formalization adds is the identification — determinism and openness as two faces of one operator, not merely compatible — and the conjunction with the no-bearer result, which removes the question’s subject as well as its dilemma. The theorems concern precisely the diagonal-bound coordinate, which is no concession, because the successor question the dissolution leaves behind is phrased in exactly those terms: which seats are self-referential, and what do their records create. The file states its own cap: no ontology of agency is claimed; what is proved is that determinism and interior openness are consistent and meet at the diagonal.

9 Five Dissolutions

Each dissolution has one anatomy: the debate’s target is a totalization — a substance, a flag, a whole, a bottom, an owner; a theorem shows the totalization is free, weightless, or resting on a presupposition that fails; and a successor question inherits everything tractable.

1. Substance dualism (mind vs. matter: which is fundamental?). Dissolved by Proposition 7.1: mind-talk and matter-talk are two vocabularies over one world class; classification is vocabulary-relative; and the correlation between the readings — the mind–body relation itself — is expressible in neither alone. The debate presupposed two candidate substances; what the theorems describe is one process and two readings, neither able to state its relation to the other. *Successor*: which architectures bind which streams, at what cost — an empirical program.

2. Consciousness attribution (is it conscious?). Dissolved by Propositions 8.2 and 7.2: the interior flag is free and its witnesses are identified by every structural vocabulary — including CTM’s and IIT’s — so no debate inside those vocabularies can terminate; and the subject term denotes a leaking factorization, so there is no invariant bearer to attribute to. *Successor*: what does this architecture integrate, and what is it blind to — decidable (Props. 7.5, 6.2) and measurable.

3. The completed totality (what is the universe, as a whole?). Dissolved by Proposition 5.1: every observable is a prefix fact; completions are invisible; “forever” is not a measurement. The whole-as-object is not wrong but weightless. *Successor*: the finite-depth structure of the ongoing process.

4. The ground (why is there something? what is at the bottom?). Split by Propositions 3.5 and 3.4: the certificate question (“am I fundamental?”) is interior-undecidable; the shape question is answered — well-founded answered demand chains terminate in a ground whose law is $x = \neg x$, given R1–R4. The debate wanted a certificate of the terminus; what exists is a theorem about the chain. *Successor*: what the one ground forces.

5. Free will (could I have done otherwise?). Dissolved by Proposition 8.3: determinism and interior openness are one operator’s two faces, and the question’s presupposed owner does not exist. The three classic camps — libertarian author, hard-determinist puppet, compatibilist desirer — share the failed presupposition. *Successor*: which seats are self-referential — where the diagonal binds — and what their records create.

10 What Stays Open

The refusals of the position are findings in their own right. Each question below is one the position does not answer, paired with the theorem showing that, in the vocabulary the position allows itself, it cannot be answered. The open status is the result, not an omission.

Question	Why it stays open
Whether the whole (or anything) is experiencing	Family.lean (family_interior_is_free, every width); Witness.lean (the registering reading constructed as conservative, axiom-free); Schema.lean (settling it requires a non-structural primitive)
Whether the ground is fundamental or derived	Aseity.lean (no consistent system certifies its own ground; axiom-free)
The completed totality (the finished whole, the continuum, the “forever” fact)	Totalization.lean (weightless at every finite depth); CompletionFree.lean
Other minds	InteriorExperience.lean (other_minds_free, mind_count_free)
Which branch; any value-class fact from the law	ValueForm.lean; SelfDerivation.lean
The cutoff scale / absolute units	PlanckFree.lean (axiom-free)
Anything beyond occurrence from the first-person datum	Cogito.lean (the cap is part of the theorem)

11 The Assumptions

A premise is a named joint between a theorem and its informal reading. The table below is the complete list of such joints: each row states how the premise could be challenged and which propositions would lose their application to the actual world — never their validity — if it failed. No proposition anywhere in the position relies on a premise absent from this table, and no L1 proposition loses validity under any premise failure.

#	Premise	How to challenge it	What depends on it
B1	The $x = \neg x$ substrate class models our physical universe	Experiment (§12)	Instance reading of Prop. 7.1; all L2/L3 faces. L1 untouched
B2	The certified describer class models minds	Neuroscience of binding	Instance reading of Prop. 7.1; the relevance of Prop. 6.2 to us. Props. 6.1, 6.3, 7.2 untouched
B3	Seat-indexed quantities read as observers/perspectives	Exhibit a non-perspectival seat-indexed quantity	Philosophical reading of Prop. 4.4; its mathematics untouched
R1	Every seeded specification's demand for a ground is answered	An unanswered demand — a brute contingency	Prop. 3.4 only
R2	No infinite regress of grounds (well-foundedness of demand)	Infinite-descent metaphysics	Prop. 3.4 only
R3	Structureless = equivariant (grounded: equivariant = definable from the distinction alone, <code>Structureless.lean</code>)	A reading of “importing no structure” outside pure-equality definability	Prop. 3.4 only
R4	An unseeded ground is generative: its law moves something	A static “from nothing”	Prop. 3.4 only
H	The law reads every seat alike (homogeneity)	An inhomogeneous law	Prop. 6.4 only
O	A finite observer's measurement is a finite-depth observable on histories (robust: = uniform continuity = finite-prefix factorization, <code>Continuity.lean</code>)	Exhibit a measurement that is not uniformly finitary	The measurement readings of Props. 5.1–5.3 and the interior half of Prop. 6.4; the mathematics untouched
F1	<code>CTMConservativity.lean</code> faithfully transcribes the published CTM	Exhibit a published CTM predicate that reads the flag	The CTM half of Prop. 7.2; the schema level (<code>Schema.lean</code>) untouched
F2	<code>IIT.lean</code> faithfully transcribes the Kleiner–Tull axiomatization	Exhibit an IIT predicate not determined by cause-effect structure	The IIT half of Prop. 7.2; the schema level untouched

One structural remark. The attack that would fell F1 or F2 — exhibiting a published CTM or IIT predicate that reads the interior flag — would thereby show that the theory in question contains a non-structural primitive, which is `Schema.lean`'s other horn (`deciding_theory_not_structural`). Either the published vocabularies are structural, and the conservativity theorems apply to them; or they are not, and they contain, by their own texts, exactly the non-structural primitive whose necessity Proposition 8.1 asserts. The premises can fall; the disjunction they serve cannot.

12 What Experiment Can Touch

The position inherits refutation conditions through its instance premises; experiments can strand the instance, never the schema, and no experiment can confirm a free proposition (Prop. 8.1). Through B1, summarized from the external physics instance quarantined behind it: a dark-fermion mass forced, with zero adjustable parameters, into a 3–5 GeV window, so that direct-detection exclusion of that window refutes B1; a dark-energy equation of state pinned at $w_0 = -0.885$ with no free dials; and Lorentz-violation searches, where any positive sidereal signal above quadratic suppression refutes B1. In each case what falls is B1, the measured faces of Propositions 4.2–4.3, and the instance reading of Proposition 7.1; every L1 proposition stands, by the layer discipline.

13 The Position in the Literature

The position belongs to two lineages, and its relation to each can be stated plainly.

The process lineage. Most of the individual theses have ancient or modern assertions, and the honest description of this work’s relation to them is: a proof where there was an assertion, or an explicit assumption where there was an implicit one. Heraclitus asserted flux and the denial of static being; Proposition 3.2 is that assertion with a model, and Proposition 3.3 adds that stasis-denial and perpetual motion are one hypothesis. Nāgārjuna denied intrinsic essence (svabhāva) and asserted dependent origination [5]; Propositions 4.3 and 5.1 are the two halves as theorems — persistence requires interaction, and the completed totality answers to no measurement. Spinoza put one substance under two attributes [1]; Proposition 7.1 puts one process under two vocabularies and adds the theorem that their correlation is storable in neither alone. Whitehead made process primary and creativity ultimate [7]; Proposition 4.5 makes novelty a theorem. Rescher systematized process over substance [19]; here every clause of such a system is tagged, checked, or proved free. Wheeler’s “it from bit” [16] sought physics in binary distinctions; Proposition 3.1 shows the distinction’s dynamics is forced rather than chosen, and §12 states what would refute the instance. Rovelli relativized facts to systems [17]; Proposition 4.4 attaches the self-derivation no-go. Ladyman and Ross eliminated things in favor of structure [23]; here the structural core is prelude-only Lean rather than argument. The Conscious Turing Machine [28] and Integrated Information Theory [29, 30] are structural theories of consciousness; Proposition 7.2 proves conservativity inside their own vocabularies, on the fidelity premises F1/F2.

The formal lineage. The second lineage is the use of formal tools on philosophical systems, and it is the one against which this paper’s central claim is measured. Spinoza’s *ordine geometrico demonstrata* [1] is the ambition the title cites: expose every definition, axiom, and inference to inspection. Its machine-age instances have so far concentrated on single arguments. Gödel’s ontological proof [20] was brought to automated and interactive verification by Benz Müller and Woltzenlogel Paleo [26], work that produced genuinely new philosophical findings — confirmation of the modal collapse, and later the inconsistency of Gödel’s original axiom set [27]; Oppenheimer and Zalta used an automated prover to simplify the ontological argument [24], within a broader program of computational metaphysics [21]. This paper

differs from those precedents in scope and in direction. In scope: what is checked is not one argument inside a position but a position — six parts and a companion ethics, thirty-five self-contained files. In direction: the earlier mechanizations checked arguments for conclusions their authors hoped to establish; a large fraction of this corpus proves instead that questions cannot be answered, and those undecidability results *are* the position’s content, not its failures. To our knowledge, no metaphysical system before this one has had every load-bearing claim kernel-checked, every assumption enumerated with its dependents, and every refusal carried by a theorem. That condition — not any single thesis — is what this paper claims as new.

What is new in content. Setting the format aside entirely, six claims of the position are, to our knowledge, not in the prior literature, and each can be located exactly by where its nearest precedent stops.

1. *The generative reading of the diagonal.* The tradition reads fixed-point results as limitations — walls around what systems can do. Here the same failure is the motor: the founding specification’s unsatisfiability is its generativity, and the Gödel-style guarded fixed point *is* the two-tick clock (Prop. 3.2, `godel_is_clock`). Spencer-Brown’s re-entrant form oscillates and was read as generating time [10, 14]; what has no precedent is the uniqueness clause (Prop. 3.1) and the frame-correspondence equivalence (Prop. 3.3): stasis-denial and perpetual motion as one hypothesis, as a biconditional.
2. *A third status for the interior flag.* The consciousness debate has run on three verdicts — true, false, meaningless. The silent-witness pair (Props. 7.3–7.4) establishes a fourth: *conservative* — provably costless to adopt, provably impossible to detect, and symmetric, withheld from the denier by the same theorem that withholds it from the asserter. The predicate half of Prop. 7.2 is kin to Chalmers’ structure-and-dynamics argument [18]; the constructed model pair, the conservativity results inside CTM’s and IIT’s own vocabularies, and the height-one closure of the registration hierarchy are not.
3. *The typed indexical.* Perry posed the essential indexical as a problem [13]; Ismael argued the self-locating facts are located rather than missing [22]. Proposition 6.4 classifies: the indexical is value-class — a record, real and law-underivable, of the wrong type to be law-content — with the self-derivation no-go attached.
4. *Novelty without indeterminism.* Bergson asserted the reality of the unforeseeably new [6] but offered no mechanism; Proposition 4.5 locates one — a write-once, law-underivable record stratum inside an exactly reversible law.
5. *The cost theorem for dissolution.* Carnap dissolved external questions [8]; Chalmers diagnosed verbal disputes [25]. The verdict calculus adds what neither had: a theorem pricing the settlement of a free question exactly (a primitive that separates the witness pairs, Prop. 8.1), a proof of interminability below that price (Prop. 8.2), and a discipline of successor questions that keeps the tractable remainder (§9).
6. *Value pinned to the record stratum.* In the companion ethics, the one normative premise attaches value to what is written — so that the priority of prevention over remedy follows from how the register works (compensation adds records; nothing erases,

NoErasure.lean) rather than from a sensibility, and moral practice is keyed to the decidable structure of a candidate patient rather than to a flag proved permanently free.

One entry is listed as descent rather than novelty, deliberately: the diagonal account of interior openness (Prop. 8.3) descends from MacKay’s logical indeterminacy [9]. What is claimed there is the identification of the two faces and the conjunction with the no-bearer result — the conjunction, not the mechanism.

What is new in format. Four features have, to our knowledge, no prior in either lineage: refusals carried by theorems (§10); premises with declared dependents (§11); a metaphysics inheriting empirical refutation conditions from a physics instance (§12); and undecidability results presented as a position’s content rather than its boundaries.

14 Objections and Replies

Objection 1: the formalization gap. *The theorems are about small discrete structures — Booleans, tick-indexed histories, finite alphabets. The grand vocabulary (“being,” “novelty,” “self”) is doing all the philosophical work, and none of it is checked.*

The gap is real, and the method’s response is to make it visible rather than to deny it. Every passage from theorem to metaphysical reading is one of exactly three things, each visible on the page. Where the reading is definitional (a seat *is* an index a quantity is read from; a boundary *is* a factorization), the definition is stated in §2 and formalized in the cited file, and a reader who rejects it rejects a definition, not an argument. Where the reading is a substantive identification, it is a named premise with declared dependents, and the table in §11 is the complete list. Where it is neither, it is tagged INTERPRETATION and severable.

The objection is also right about distribution, and the honest summary should own it: the corpus is not uniform in mathematical weight. In §3 the mathematics is deliberately small — the classification behind Proposition 3.1 is an exercise, and the metaphysical weight sits openly on R1–R4. In §§6–7 the situation inverts: Family.lean proves class membership, rigidity, and conservativity kernel-only at every ring width, and CTMConservativity.lean carries a full transcription of a published theory. The forced core is smallest where the theses are oldest — Heraclitus did not need help — and largest where they are newest. What machine-checking buys overall is not that the metaphysics is proved; it is that the ratio of proved to assumed is exact and public, and that the assumed part is enumerable in one table.

Objection 2: proving undecidability is not neutrality. *Calling the interior flag “free” quietly favors the deflationist: the position refuses to say anything is conscious, which in practice reads as denial.*

The position is symmetric where the theorem is, and the symmetry is constructed, not asserted. The silent-witness pair (Props. 7.3–7.4) builds *both* models — a world in which the process registers itself at every tick, and one in which it does not — and proves them observation-identical: adopting the registering reading is provably costless, and denying it is provably unfounded. Meanwhile occurrence itself (Prop. 6.3) is affirmed as strongly as anything in the system: seated, necessary in its validity, self-refuting to deny. What is withheld is exactly the attribution — the flag pinned to a bearer — and it is withheld from the denier on

the same theorem that withholds it from the asserter. A deflationism that concluded “therefore nothing is experiencing” would be refuted by `Witness.lean` as directly as its opposite.

15 Conclusion

The position, compressed to its architecture: one forced shape at the ground; a forced order of time; a write-once, law-underivable stratum of records in which novelty consists; a budget of observation that renders completed totalities weightless; seats without selves and occurrence without owners; one process under two vocabularies, with the mind–body relation statable in neither alone; a verdict calculus that proves certain debates interminable inside their vocabularies; and a free-will antinomy dissolved into two faces of one operator. Around the theorems, a discipline: five verdicts covering every clause, eleven premises covering every joint, an exposure table stating what experiment can reach, and refusals that are themselves theorems.

A companion document in the same repository — *Ethica, ordine mechanico demonstrata* — writes ethics to the same standard. Its structural results are carried by the same corpus. Hume’s is/ought gap [3] arrives as a theorem about vocabularies, together with its exact converse: any vocabulary that expresses a normative claim contains a primitive beyond the descriptive one, so every ethics either takes the supervenience horn or carries at least one bridge premise — the gap becomes the theorem that every ethics has a premise table. The moral-status, desert, and aggregation impasses dissolve on the witness, stream, and horizon theorems respectively. The document carries exactly one normative premise, tabled with its dependents, and closes by proving that it cannot prove its own bridge (`settling_needs_new_primitive`, applied to itself). Four proof files of its own — `Dominance.lean`, `Bridge.lean`, `NoErasure.lean`, `Comparison.lean` — carry what it needed beyond the treatise’s corpus.

Nothing above asks to be believed on the strength of its prose. The forced part checks — in minutes, with a bare `lean` binary; the free part is proved free; and the assumed part is eleven rows long.

References

- [1] B. Spinoza. *Ethica, ordine geometrico demonstrata*. Amsterdam, 1677.
- [2] R. Descartes. *Meditationes de prima philosophia*. Paris, 1641.
- [3] D. Hume. *A Treatise of Human Nature*. London, 1739–40.
- [4] I. Kant. *Kritik der reinen Vernunft*. Riga, 1781; 2nd ed. 1787.
- [5] Nāgārjuna. *Mūlamadhyamakakārikā* (The Fundamental Verses on the Middle Way). c. 2nd century CE.
- [6] H. Bergson. *L’Évolution créatrice*. Alcan, Paris, 1907.
- [7] A. N. Whitehead. *Process and Reality*. Macmillan, New York, 1929.
- [8] R. Carnap. Empiricism, semantics, and ontology. *Revue Internationale de Philosophie* 4:20–40, 1950.

- [9] D. M. MacKay. On the logical indeterminacy of a free choice. *Mind* 69(273):31–40, 1960.
- [10] G. Spencer-Brown. *Laws of Form*. Allen & Unwin, London, 1969.
- [11] F. W. Lawvere. Diagonal arguments and cartesian closed categories. In *Category Theory, Homology Theory and their Applications II*, Lecture Notes in Mathematics 92, Springer, 1969.
- [12] J. Hintikka. Cogito, ergo sum: inference or performance? *The Philosophical Review* 71(1):3–32, 1962.
- [13] J. Perry. The problem of the essential indexical. *Noûs* 13(1):3–21, 1979.
- [14] L. H. Kauffman and F. J. Varela. Form dynamics. *Journal of Social and Biological Structures* 3(2):171–206, 1980.
- [15] D. Parfit. *Reasons and Persons*. Oxford University Press, 1984.
- [16] J. A. Wheeler. Information, physics, quantum: the search for links. In W. Zurek (ed.), *Complexity, Entropy, and the Physics of Information*. Addison-Wesley, 1990.
- [17] C. Rovelli. Relational quantum mechanics. *International Journal of Theoretical Physics* 35:1637–1678, 1996.
- [18] D. J. Chalmers. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press, 1996.
- [19] N. Rescher. *Process Metaphysics: An Introduction to Process Philosophy*. SUNY Press, Albany, 1996.
- [20] K. Gödel. Ontological proof. In *Collected Works, Volume III: Unpublished Essays and Lectures*. Oxford University Press, 1995.
- [21] B. Fitelson and E. Zalta. Steps toward a computational metaphysics. *Journal of Philosophical Logic* 36(2):227–247, 2007.
- [22] J. Ismael. *The Situated Self*. Oxford University Press, 2007.
- [23] J. Ladyman and D. Ross. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press, 2007.
- [24] P. Oppenheimer and E. Zalta. A computationally-discovered simplification of the ontological argument. *Australasian Journal of Philosophy* 89(2):333–349, 2011.
- [25] D. J. Chalmers. Verbal disputes. *Philosophical Review* 120(4):515–566, 2011.
- [26] C. Benz Müller and B. Woltzenlogel Paleo. Automating Gödel’s ontological proof of God’s existence with higher-order automated theorem provers. In *ECAI 2014*, IOS Press, 2014.
- [27] C. Benz Müller and B. Woltzenlogel Paleo. The inconsistency in Gödel’s ontological argument: a success story for AI in metaphysics. In *IJCAI 2016*, 2016.
- [28] L. Blum and M. Blum. A theory of consciousness from a theoretical computer science perspective: insights from the Conscious Turing Machine. *Proceedings of the National Academy of Sciences* 119(21):e2115934119, 2022.

- [29] G. Tononi, M. Boly, M. Massimini, and C. Koch. Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience* 17:450–461, 2016.
- [30] J. Kleiner and S. Tull. The mathematical structure of integrated information theory. *Frontiers in Applied Mathematics and Statistics* 6:602973, 2021.
- [31] L. de Moura and S. Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, LNCS 12699, Springer, 2021.