

The Minimum Distinction as Fundamental Unit: Regress, Structure, and Duality

Keith Taylor versf-eos.com

General Reader Summary

What is the smallest thing reality is made of? This paper argues there is a clean answer: it must be the *minimum distinction* — the smallest possible difference between one thing and another — and nothing else.

The core argument runs by elimination. Anything larger than the minimum distinction (atoms, particles, fields, spacetime points, abstract structures) can always be broken down into smaller distinctions among its parts; those smaller distinctions are more fundamental than the larger object they compose. Keep breaking down, and you must eventually hit something that cannot be broken down further. That something is, by definition, the minimum distinction. The only alternative is that there is no floor at all — that distinctions keep getting smaller forever without end — but such a reality could not support stable facts, records, or measurements, so it would not be a reality compatible with physics as we have it.

The paper then characterizes what the minimum distinction must look like: it has two sides (any richer structure decomposes into two-sided pieces), a boundary between them (otherwise there is no place where the distinction holds), and a direction (one side is committed, the other is not — otherwise the distinction supports no record). Described geometrically this object is called a *fold*; described informationally it is a *bit*. These are two names for the same structure — the geometric and informational faces of the most basic unit of actualized reality.

Four defensive moves support the argument. *First*, the reasoning is shown to survive against every major contemporary candidate for a more fundamental primitive, including those proposed in physics research programmes (quantum fields, entanglement, spin networks, causal sets, strings, constructor theory) and in philosophy (relations-first ontologies). *Second*, the fold is proved — not merely argued — to be the unique primitive satisfying four named admissibility conditions: minimality, boundary-realization, orientation, and composability into the next structural level (the minimum fact, established elsewhere to require exactly seven connecting elements, or *K = 7 closure*). This uniqueness is the paper's main theorem, derived constructively from the four conditions in four steps. *Third*, the fold is forced independently by two separate routes: once from informational considerations (regression to the minimum distinction) and once from geometric considerations (the boundary structure the Void requires). These routes converge because any admissible primitive must satisfy both constraints simultaneously, and the fold is precisely the object in their intersection. *Fourth*, the framework is anchored to a derivational track record: the fold/bit primitive underpins published derivations of the fine-structure constant

(to 15 parts per million), the cosmological constant, the Standard Model particle spectrum, and several other observable quantities of physics.

The paper is explicit about what remains open. A more abstract mathematical reformulation of the uniqueness theorem (in category-theoretic terms — a refinement of how the result is expressed rather than a gap in the result itself), a full formal treatment of how fold-and-bit descriptions combine at every scale of composition, and a rigorous alternative formulation in which relations rather than things are taken as fundamental are all flagged as targets for follow-up work. What this paper establishes is the foundational step: the fundamental unit of reality is the minimum distinction, it has a specific structure that can be characterized precisely, and that structure is derived as a provable necessity under explicit admissibility conditions — not an arbitrary choice, not a stipulation, not a selection among alternatives.

Contents

- General Reader Summary
- Abstract
- 1. Why the unit matters
- 2. Why the fundamental unit must be the minimum: the regress argument
 - 2.1 The argument
 - 2.2 What the argument rules out
 - 2.3 Where the regress terminates
 - 2.4 What the argument does not establish
 - 2.5 Holism, fields, and relational ontologies
 - 2.6 Specific research programmes
- 3. The minimum distinction is a minimum asymmetry
 - 3.1 Why binary
 - 3.2 Why a boundary
 - 3.3 Why orientation
 - 3.4 Summary
- 4. The geometric aspect: the fold
 - 4.1 Abstract universality at the unit level
 - 4.2 Necessity criteria for the fold at the composition level
 - 4.3 Uniqueness of the fold under dual constraints
 - 4.4 Pre-metric status
- 5. The informational aspect: the bit
- 6. Fold and bit as dual descriptions
 - 6.1 The mapping at the unit level, and where the real duality lives
 - 6.2 Implications of the duality
- 7. What is built from the minimum: distinction, fact, record
 - 7.1 Two foundational levels, not in competition
- 8. Location of the unit relative to observables
- 9. Claimed, assumed, open
- Conclusion
- Appendix A: Anticipated Questions and Responses
- References

Abstract

The layered foundation of VERSF identifies two co-primitive aspects — potentiality (the Void) and actualization (distinguishability / commitment) — without yet fixing a *unit* for the programme to build on. This paper fixes the unit in two steps. First, a regress argument shows that the fundamental unit of actualized reality must be the minimum distinction, and nothing else: any larger candidate contains distinctions as parts, which are structurally prior to it. The argument is defended against the principal holistic objections (field-primary, entanglement-primary, relations-first) by showing that each of these still operates through distinctions at the level of its content; relational-first ontologies face a trilemma from which no current formulation we are aware of escapes. The argument is then tested against specific live research programmes — Wheeler's "it from bit" (compatible precursor), algebraic QFT, loop quantum gravity, causal set theory, string theory, and constructor theory — each of which identifies a fundamental structure whose content is, at its own level, delivered through distinctions. The regress places each of these programmes correctly above rather than below the minimum distinction. This rules out composed folds, $K = 7$ closures as *fundamental units* (though $K = 7$ is retained as the minimum closure at a higher structural level), extended structures, continuous fields, and higher-dimensional objects. Second, the minimum distinction is characterized structurally: it has two sides (binary, by decomposition), a boundary (by separability), and an orientation (by the requirement that commitment support a record, tightened by appeal to symmetry \rightarrow indistinguishability and to the orientation–irreversibility link). At the unit level, the fold and other binary-oriented primitives (graph edge, asymmetric binary relation, oriented pair) are abstractly isomorphic; what distinguishes the fold is its substrate realization, which is selected by *necessity criteria* — not preference — on composition into $K = 7$ closure. This selection is formalized in §4.3 as a uniqueness theorem: any primitive satisfying the admissibility constraints (C1) minimality, (C2) boundary realizability, (C3) orientation, and (C4) closure admissibility is isomorphic, up to relabelling of substrate regions, to a fold. The theorem transforms what §4.2 establishes by exhaustion over canonical candidates into a formal uniqueness statement under a named constraint set, with the two-route convergence (regress + void boundary) following as a corollary. The same fold object is derived by an independent topological route in the companion "Fold Interface Law" paper; the two-route convergence is not accidental but structural — the fold is the object in the intersection of the void's boundary-forcing constraint and the regress's minimal-asymmetry constraint. The composition-level duality between fold and bit descriptions is stated as a formal target ($\tilde{\varphi} : \mathcal{F} \rightarrow \mathcal{B}$ with agreeing and opposing composition operations correspondent) with full proof at arbitrary order deferred to a companion technical paper. The fold's binary-plus-orientation structure maps explicitly to \mathbb{C}^4 under the minimal reversible representation, closing the first link of the $\mathbb{C}^4 \rightarrow \mathbb{CP}^3$ geometric chain used in later programme papers. The foundational argument is anchored to a derivational track record: the fold/bit primitive underpins the programme's derivations of α^{-1} to 15 ppm, $K = 7$ from six convergent routes, the cosmological constant, the Standard Model spectrum from hexagonal geometry, and the κ -field mass coefficient. The paper closes by marking what is forced, what is imported, and what remains open.

1. Why the unit matters

A layered foundation — potentiality and actualization as co-primitives — is not yet a physics. To do work, the foundation needs a *unit*: a characterization of the smallest thing that actualized reality is made of, from which all further structure accumulates.

Two questions follow. *First*: what *must* the fundamental unit be? The programme needs a reason why one candidate rather than another is fundamental, not a stipulation dressed as a choice. *Second*: what does that unit *look like*, structurally? Once the unit is fixed in principle, its features have to be characterized precisely enough for the rest of the programme to build on them.

The first question is addressed in §2 by a regress argument, defended (§2.5) against the principal holistic objections and tested (§2.6) against specific live research programmes (Wheeler, AQFT, LQG, causal sets, strings, constructor theory): the fundamental unit must be the *minimum distinction*, because every larger candidate fails by its own structure, and no currently active programme identifies a more fundamental primitive that escapes the regress. The second question is addressed in §§3–6: the minimum distinction has two sides, a boundary, and an orientation, manifest as fold (geometrically) and bit (informationally), which are dual descriptions of one structure. §4.3 states the central uniqueness result as a formal theorem: any primitive satisfying four explicit admissibility constraints is isomorphic to a fold up to substrate relabelling. §§7–8 locate this unit within the programme's hierarchy, note the bridge to observables, and anchor the foundational argument to the programme's derivational track record. §9 marks what is established, what is imported, and what is open.

A note on vocabulary. "First" and "minimum" in this paper are logical-structural, not temporal — time has not yet emerged at the layer discussed here. "Minimum distinction" means the least structurally complex actualization, not the chronologically earliest one.

The programme uses several equivalent terms for this primitive, depending on vocabulary context: *minimum distinction* and *minimum asymmetry* (structural register), *minimal distinguishability* (distinguishability-theoretic register), *fold* (geometric register), *structural bit with direction* (informational register). These are the same object described in different vocabularies, not successive refinements of different concepts. The paper moves freely between them as the argument requires.

2. Why the fundamental unit must be the minimum: the regress argument

2.1 The argument

The argument is a regress-blocker. It rests on three premises.

Premise 1. The programme is committed to distinctions as the content of actualized reality. Whatever is actual is actualized as, in, or by distinctions. (Imported from the layered foundation and not argued here.)

Premise 2. If a candidate fundamental unit U has internal structure — if U contains distinctions as parts — then those parts are themselves distinctions, and they are *structurally prior* to U (in the sense of being decomposition-constituents of U , not in any spatial or metric sense — spatial extent has not yet been introduced at this layer).

Premise 3. Parts are prior to wholes composed of them. A thing made of distinctions is not more fundamental than the distinctions it is made of; the distinctions are what it is made of.

From these: if U has internal structure, then U is not fundamental. Its decomposition-prior parts are more fundamental. The same argument applies to U 's parts. If any of them has internal structure, they too fail to be fundamental. The regress terminates only when a candidate is reached that has no internal distinctions — a unit with no decomposition-prior distinctions as parts.

A unit with no internal distinctions is, by definition, the minimum distinction. The regress terminates there because below the minimum distinction lies the Void, which is not composed of distinctions and is therefore not a smaller candidate unit at all. The floor is structurally determined: it is where distinctions end.

Conclusion. The fundamental unit must be the minimum distinction. Any larger candidate invites the question "why not more fundamental?", and the question is answered — "because below this there are no distinctions, only potentiality" — only at the minimum.

2.2 What the argument rules out

The regress does not merely establish the minimum distinction as a possible answer. It rules out every alternative. Working through the main candidates:

- *Composed folds or composed bits.* If the fundamental unit were a fold composed of sub-folds, the sub-folds would be prior. Ruled out.
- *$K = 7$ closures.* The $K = 7$ closure is the minimum *fact* — the minimum stable, closed structure that supports a record — and it is foundational at its own level. But it contains distinctions as parts, and those parts are prior. The $K = 7$ closure is therefore the minimum fact but not the fundamental unit. Ruled out as fundamental unit; retained in its proper role as minimum closure. (See §7.1 for the full framing of the two-level architecture.)
- *Extended structures generally.* Any structure with spatial, temporal, or relational extension has, in that extension, distinctions between its parts — "here" versus "there," "now" versus "then." These are prior to the extended structure. Ruled out.
- *Continuous fields and smooth manifolds.* A continuous field has distinguishable values at distinguishable points; a manifold has distinguishable locations. Each such distinguishability is a prior distinction. Continuous structures are, at best, large-scale approximations to accumulated minimum distinctions. Ruled out.
- *Strings, loops, branes, higher-dimensional objects.* These are extended structures with internal distinctions (points on the string, positions on the brane). Ruled out by the extension argument.

- *Mathematical structures treated as physical primitives — groups, algebras, categories.* These are structures of relations among distinct elements. The elements are distinctions, and they are prior to the structures they compose. Fundamental in mathematics, not fundamental as physical units. Ruled out.

The survivor is the minimum distinction, and nothing else.

2.3 Where the regress terminates

The regress halts at a unique structural point. It does not halt at an arbitrary threshold chosen by convention. It halts where distinctions end — where further descent would leave the domain of the actualized altogether and enter the Void.

Below the minimum distinction is not a more fundamental minimum distinction. Below the minimum distinction is the Void: pure potentiality, undifferentiated, containing no actualized distinctions. The Void is not a more fundamental candidate unit; it is not a distinction at all. Going "below" the minimum distinction does not mean finding a more fundamental distinction; it means leaving the domain in which "a distinction" is the right kind of thing to look for.

This is also why the answer is unique. Any putative alternative fundamental unit either contains distinctions as parts (and is ruled out by the regress) or is not itself a distinction (in which case it is the Void, not a unit). The two-option exhaustion forces uniqueness — conditional on the decomposition relation being well-founded, i.e., on there being no infinite descending chain of distinctions; this assumption is flagged explicitly in §9.

2.4 What the argument does not establish

The regress is powerful but narrow. It is worth being explicit about the load it bears and the load it does not.

- It does not establish *what the minimum distinction looks like*. The argument establishes that there must be a minimum distinction (given well-foundedness) and that the fundamental unit must be it. The structural characterization — two sides, boundary, orientation, manifest as fold and bit — is the work of §§3–6.
- It does not establish that the minimum is *irreversibly committed*. Whether the minimum asymmetry is a stable commitment or a reversible fluctuation is a further question. The commitment principle of the broader programme settles this; the regress argument is neutral.
- It does not establish the *existence* of anything actualized. The argument is compatible with no actualization occurring at all: it characterizes the structure of actualization *conditional on* actualization occurring, and is silent on whether actualization occurs.

These qualifications clarify where the argument's load is. It is doing foundational structural work — uniqueness of the fundamental unit — not metaphysical work (whether anything exists) and not descriptive work (what the unit looks like). §2.5 addresses the principal objection to the argument itself, and the rest of the paper takes up the descriptive task.

2.5 Holism, fields, and relational ontologies

Premise 3 — parts are prior to wholes — is the assumption most likely to be challenged, and the challenge comes from physics as much as from metaphysics. Field-theoretic ontologies treat fields as fundamental; entanglement is often read as evidence that wholes are primary over their apparent parts; relational ontologies claim that relations are prior to the relata they relate. If any of these holds strictly, the regress does not apply in the simple form given.

The reply, in short, is that each of these alternatives, on inspection, still operates through distinctions — the distinctions are simply located in a different structural register than "spatial parts of wholes." The regress runs on decomposition of content, not on spatial mereology, and survives translation into each of the alternative ontologies.

- *Field-theoretic ontology.* A field is "whole" in the sense that it is not decomposable into spatially localized parts. A field-primacy partisan might reply that the distinctions among modes, values, or configurations are epistemic — facts about how we describe the field — rather than ontic facts about what the field is. The reply to the reply is structural: the *content* of what a field ontology posits is individuated by distinctions. If two field configurations differ, the difference is specified through distinguishable features of them; if they do not differ in any distinguishable feature, they are the same configuration — by the identity of indiscernibles applied to the ontology's own commitments. A field ontology without decomposable content is an ontology committed to nothing at all. The regress applies to the decomposable content the ontology commits to, not to the epistemic vocabulary used to describe it. The field's "wholeness" is therefore orthogonal to the argument: the field is not a primitive that bypasses distinctions; it is a substrate whose content is specified by distinctions.
- *Entanglement holism.* Entangled states are non-separable in the sense that the whole state is not the product of its subsystem states. But the formalism of entanglement is built on distinguishing state vectors in Hilbert space, distinguishing measurement outcomes, and distinguishing correlations from independence. An entangled state is not a "no-distinction" state; it is a state whose distinctions live in a different tensor structure than the separable case. Holism about separability does not imply holism about distinction-content. The regress applies to the distinction-content, which is present in entangled and separable states alike.
- *Relational ontology.* Relations are claimed to be prior to relata. Any such programme faces a trilemma. *Horn (a):* the content of the primitive relations is specified by what they relate, in which case the relata are structurally prior and the regress applies normally. *Horn (b):* the content is specified without reference to relata at all, in which case the relations are underdetermined — there is no fact of the matter about what the relations relate — and the ontology lacks theoretical content. *Horn (c):* the content is specified by reference to further relations (meta-relations), opening a second-order regress in which the meta-relations must themselves be specified by content; this regress either terminates at some minimum relational structure, whose content is then fixed by whatever it relates (returning to horn a), or runs indefinitely (violating the well-foundedness that any meaningful theory requires). No current formulation we are aware of escapes all three horns. The programme's position is that distinguishability is itself a relation that

presupposes its relata, placing it on horn (a) and making the relata structurally prior in the sense Premise 3 requires. A fully-worked escape from the trilemma remains a genuine open question, flagged in §9; but the burden is correctly located on the relational-first programme, not on this one.

The broader point is that Premise 3 is not about spatial parts of spatial wholes. It is about decomposition-constituents of structural content. Even ontologies that deny spatial mereology accept that their theoretical content has decomposable structure, and the regress runs on that structure. A referee who wants to defeat the regress by appeal to holism or field-primacy would need to show that the holistic ontology has *no decomposable content at all* — no distinct modes, no distinct measurement outcomes, no distinct relational patterns. That is a much stronger claim than the usual holism, and one that no current physics or metaphysics defends.

2.6 Specific research programmes

§2.5 addresses the general shape of the alternatives. It is worth checking the defense against specific research programmes that make concrete fundamental-unit claims, because a referee will want to see the argument survive encounter with the live physics literature, not just with generic positions. Six programmes are worth naming.

- *Wheeler's "it from bit"*. Wheeler's proposal that information is prior to matter is, on the reading adopted here, a *precursor* to the fold/bit framework, not a rival. Wheeler identifies the bit as fundamental; the present paper identifies the minimum distinction (with fold and bit as dual descriptions) as fundamental. The two positions agree that information-theoretic primitives are foundational; the present paper adds the structural argument (regress) that forces this, the derivational characterization of the minimum, and the composition-level duality with fold geometry. Wheeler's slogan survives as the informational face of the programme. Named here as compatible, not contested.
- *Algebraic Quantum Field Theory*. AQFT treats nets of operator algebras as fundamental structure, with local observables as primary. But operators are individuated by their action on states, and operator algebras are built from operators distinguished by their spectra, commutation relations, and effects. The algebra is not a single undifferentiated whole; it is a structure of distinguished elements with distinguished relations. The regress applies to that structure at its unit level. AQFT is therefore foundational at its own (higher) level but not at the level the regress reaches.
- *Loop Quantum Gravity*. LQG takes spin networks as primitive — graphs with nodes and edges labelled by representation-theoretic data. Nodes are discrete; edges are discrete; labels are distinguished. Every element of a spin network is a distinction. A spin network is therefore a *composed* structure in the regress sense, and the regress applies to its constituent distinctions. LQG is consistent with the fold/bit framework at a higher structural level — spin networks being specific composed configurations of the underlying fold/bit primitives — but is not more fundamental than the fold/bit.
- *Causal set theory*. Causal sets take events and causal relations (precedence) as primitive. Events are distinguished; the causal relation distinguishes predecessors from successors; both are distinction-content. The regress applies at the level of the distinguishing structure itself. Causal sets are a specific realization of oriented distinction-relations —

arguably a macroscopic realization of accumulated fold composition with consistent orientation — but the primitives of the causal set formalism are not more fundamental than the minimum distinction.

- *String theory*. Strings are extended objects with internal distinctions (points on the string, vibrational modes, positions on a brane). They are ruled out as fundamental units by the extension argument of §2.2: any extended object decomposes into prior distinctions. String theory is a higher-level theory that may be derivable from more fundamental structure, in the sense §8 gestures at for continuous-field programmes.
- *Constructor theory*. Constructor theory takes the possibility and impossibility of physical transformations (tasks, counterfactuals) as primitive. This is the most interesting case among the rivals, because tasks are not obviously composed of distinctions in the way spatial wholes are. But a task is individuated by its inputs, outputs, and the transformation between them — each of which is a distinguishable state or process. A task presupposes distinguishable states and distinguishable transformations, so the content of constructor-theoretic primitives is delivered through distinctions. The regress applies at the level of that content. Constructor theory offers a genuinely different vocabulary for physics and is compatible with the fold/bit framework at its own level, but its primitives are not more fundamental than the minimum distinction.

Across these six programmes, the pattern is the same: each identifies a fundamental structure for its own purposes, and each of those structures is composed, at some level, of distinctions. The regress applies at the distinction level. No currently active fundamental-physics programme identifies a primitive that is not itself a structure of distinctions, so the fold/bit framework does not conflict with any of them at their own level — it places them correctly relative to a more fundamental primitive that the regress forces.

3. The minimum distinction is a minimum asymmetry

The Void is pure potentiality. Structurally, this means it is perfectly symmetric: no direction is privileged, no configuration selected, no possibility fixed. The first departure from pure potentiality is therefore a single asymmetry where before there was none.

The minimum asymmetry has a definite structure. Two of its features follow from what "a distinction" means at all; the third is forced by what makes an actualized distinction capable of supporting a record.

3.1 Why binary

Claim. The minimum distinction is binary: it has exactly two sides.

Argument. Consider any candidate non-binary primitive — say, a ternary distinction with three regions A, B, C. To specify this distinction is to specify three separations: A from (BUC), B from (AUC), C from (AUB). Each separation is itself a binary distinction. Specifying the consistency among them — that A and B are non-overlapping, that boundaries meet cleanly, that $A \cup B \cup C$ is exhaustive of the partitioned substrate — requires further binary distinctions: overlap versus no-overlap, boundary-meeting versus boundary-mismatch, exhaustive versus remainder. The

ternary primitive thus decomposes into a bundle of binary distinctions (the three separations plus the consistency binaries); it is not more primitive than the binary distinction, it is a composite of several binaries.

The general form: *distinguishing* is inherently the operation of dividing "this" from "not-this." Any partition into $n \geq 3$ regions admits decomposition into a sequence of binary partitions plus binary consistency conditions, and the n-ary description adds only the bundling of these binaries. A non-decomposable distinction with more than two sides would require a primitive relation richer than "this vs. not-this," and such a relation already presupposes distinctions to define its arity.

3.2 Why a boundary

Claim. The minimum distinction requires a boundary between its two sides.

Argument. A distinction of regions must specify (a) what is distinguished, (b) what it is distinguished from, and (c) the locus along which the distinction holds. Without (c), (a) and (b) have no point of separation: they are not two regions but at best a gradient in a single undifferentiated manifold, and "distinguished from" has no operational content — there is no answer to "where does the distinction hold?"

One might attempt a boundary-free distinction by appeal to a relational property alone (e.g., "A is hotter than B" without asserting where A ends and B begins). But such a relation presupposes that A and B are already separate regions, and their separateness is precisely what a boundary supplies. Relational-only distinctions are parasitic on the boundary structure; they do not replace it.

3.3 Why orientation

Claim. The minimum *actualized* distinction is oriented: its two sides are not interchangeable.

Argument. Orientation does not follow from "distinction" in the abstract — two distinguishable but unordered sides would already qualify as a distinction. It follows from the requirement that the distinction be *committed*, i.e., capable of supporting a persistent record.

Consider a candidate non-oriented minimum: a distinction whose two sides are symmetric, in the sense that any description of the distinction is equally valid under side-swap. Such a distinction cannot support a record. A record, by definition, is a stable entity whose content says which side was committed. If the two sides are interchangeable, there is no fact of the matter about which was selected — the record content is identical under the swap, so the distinction has not been resolved, so no record has been produced. A non-oriented distinction is therefore indistinguishable from an un-committed superposition of the two possibilities.

The argument can be tightened against the natural "relational encoding" objection — the claim that a symmetric system could still encode information through its relational or correlational structure — by making two structural links explicit.

Symmetry implies indistinguishability at the level of the symmetry. Information, by any serviceable definition, is discrimination between alternatives. Discrimination requires distinguishability. If the two sides of a distinction are perfectly symmetric under side-swap, then any description invariant under that symmetry — relational, substrate-based, or informational — is silent on what the symmetry relates, because that is what it means for the description to be invariant. A "relational encoding" that claims to extract information from symmetric elements must itself be asymmetric somewhere: the asymmetry then lives in the correlation pattern, or in the relational structure, or in the temporal ordering, and orientation returns at that level. Symmetry cannot be pushed indefinitely down the chain without eventually breaking, on pain of producing a configuration in which no discrimination is possible at all. So orientation is not avoidable by relocating into relations; it is a structural requirement of any configuration that supports non-trivial informational content.

Orientation is the structural correlate of irreversibility. Commitment is defined in the programme as *irreversible* selection; irreversibility requires a direction of change. The reason is straightforward: a reversible operation leaves before-state and after-state structurally symmetric, so no fact-of-the-matter distinguishes them; irreversibility *is* the asymmetry between before and after, and that asymmetry is the direction. Direction, in turn, is an instance of orientation — the singling out of one of the two possibilities (before / after, committed / uncommitted) as the actualized one. A reversible "commitment" preserves symmetry between committed and uncommitted states, and between before-commitment and after-commitment, which is indistinguishable from no commitment at all, and supports no accumulation of entropy or memory. Orientation is therefore tied directly to the irreversibility that makes commitment produce records that track entropy, rather than reversible operations that produce nothing retained. The link between orientation at the minimum unit and entropy increase at macroscopic scales is developed in the programme's thermodynamic papers; for present purposes, it is noted so the orientation argument is anchored to irreversibility rather than resting on "because records" alone.

Put together: commitment-as-record-producing requires orientation because symmetry kills discrimination (no information), and irreversibility requires directional asymmetry (no persistent state change without it). Orientation is not an extra stipulation layered on top of commitment; it is what commitment-as-record-producing requires in order to be a coherent operation. (The characterization of commitment as record-producing is imported from the layered foundation companion paper; the orientation argument is conditional on that characterization.)

3.4 Summary

The minimum distinction has two sides (§3.1), a boundary (§3.2), and an orientation (§3.3). The first two are entailed by what a distinction is. The third is entailed by what makes a distinction actualizable as a record, with the argument tightened by appeal to symmetry-kills-discrimination and to the orientation-irreversibility link. Together, these three features exhaust the structural content of a single committed asymmetry in the Void.

4. The geometric aspect: the fold

Seen geometrically, the minimum asymmetry is what this programme calls a **fold**.

The term needs pinning down, because several topological objects can be loosely described as "two-sided with a boundary": half-planes, oriented disks, crease patterns, bordered manifolds. These are not all the same, and which one the minimum unit actually is matters for everything built on top. §4.2 evaluates all five canonical candidates (half-plane, oriented disk, crease, bordered manifold, fold) against explicit necessity criteria.

Definition. A *fold* is the identification of two substrate regions along a shared boundary, with orientation inherited from commitment — i.e., with one of the identified regions marked as the committed side. A fold is not a crease in a pre-existing sheet (there is no pre-existing sheet); it is the minimal substrate structure that realizes, as geometry, the two-sided oriented boundary of §3.

The correspondence with the three features is direct: the fold's two identified regions \leftrightarrow the two sides of the asymmetry; the shared boundary \leftrightarrow the boundary between them; the marked region \leftrightarrow the orientation.

Convergence with the Fold Interface Law. The fold as defined here appears by two independent derivation routes in the programme, and it is worth stating the convergence explicitly so it is not mistaken for a contradiction. This paper derives the fold from the structural requirements of minimum distinction (binary + boundary + orientation), proceeding through the abstract content of a minimum asymmetry and the composition requirements at $K = 7$ closure — an informational/logical route. The companion "Fold Interface Law" paper derives the same object from topological considerations: the boundary structure forced on the void by the commitment constraint, proceeding through geometric rather than informational requirements.

These are not two constructions of different objects; they are two routes to the same object. Moreover, the convergence is not accidental. The void constraint forces boundary structure on any actualized distinction (the geometric necessity); the regress forces minimal asymmetry as the fundamental unit (the structural necessity). Any admissible primitive must satisfy both conditions simultaneously, and the fold is precisely the object in the intersection — the minimum asymmetry that carries the boundary structure the void constraint requires. The coincidence of the two derivations is therefore not a fortunate overlap but a structural consequence: the fold is forced by the intersection of geometric and informational necessity, not by either route alone. Parallel two-route derivations of α , $K = 7$, and the cosmological constant are established in their respective papers of the programme; the fold is the foundation-level instance of this methodological pattern, in which a structurally correct primitive is confirmed — and here, pinned to uniqueness — by the intersection of independent derivations.

4.1 Abstract universality at the unit level

A clarification before the selection argument, because it prevents a whole class of objection.

At the unit level, the "fold" description is not privileged over other candidate descriptions of a binary-oriented primitive. Several alternatives — the directed graph edge $\langle v_1 \rightarrow v_2 \rangle$, the asymmetric binary relation $R(a,b)$ (with $R(a,b) \neq R(b,a)$), the abstract oriented pair — all have

the same structural content as the fold at the level of a single unit: two relata, a binary relation between them, and an orientation distinguishing which is which. These descriptions are *abstractly isomorphic* to the fold at the unit level. The programme's claim is not that "fold" is the unique description of the unit's abstract structure.

A clarification to preempt an objection. The isomorphism is at the level of *unit structural form* — two relata plus an asymmetric binary relation. But these alternative descriptions come with different natural composition operations: directed graphs compose by vertex identification, asymmetric relations by relational product, abstract oriented pairs by sequencing, folds by boundary identification. At the unit level these compositional structures have no work to do — the unit level is prior to composition — but the choice of which composition operation is correct matters at every level above. That is why §4.2 is required to select the substrate realization whose native composition admits $K = 7$ closure.

What distinguishes the fold from these alternatives is its *substrate realization*. The fold's definition specifies that the two sides are *substrate regions identified along a shared boundary*, not vertices of an abstract graph or relata of an abstract relation. This is a commitment about how the primitive sits in the substrate — not about what the primitive is, abstractly. And this substrate-realization commitment is what gets selected by the composition requirements of §4.2: only the substrate-region realization composes into $K = 7$ closure in the way the programme requires.

So the paper's claim is structured in two layers:

- *At the unit level*: any binary-with-orientation primitive is adequate. Fold, directed edge, asymmetric relation, and oriented pair are interchangeable.
- *At the composition level*: only the substrate-realized fold composes into $K = 7$ closure. The alternative abstract realizations are fine as mathematical descriptions of the unit-level content but do not realize in the substrate the way composition demands.

This addresses the "why exactly this structure" objection directly. At the unit level, the answer is: the structure is binary-with-orientation, and the specific vocabulary is a matter of choice. At the composition level, the answer is: only the fold's substrate realization admits $K = 7$ closure, and this is argued in §4.2.

4.2 Necessity criteria for the fold at the composition level

Necessity criteria. These are not criteria chosen after the fact to fit the fold. They are the structural requirements implicit in the admissibility axioms of the companion $K = 7$ no-go theorem (which establishes $K = 7$ as the minimum simplicial order at which a relational substrate admits fact-production). The criteria (i)–(iii) below are a distillation of those axioms in a form directly applicable to candidate substrate realizations; an explicit derivation of the distillation from the no-go theorem's axioms is deferred to a companion technical paper. What matters for the present argument is that any primitive missing any one of them cannot, by structural necessity, participate in $K = 7$ closure. The criteria are:

(i) pairwise sharing of codimension-1 boundary elements (i.e., boundary elements one dimension lower than the regions they separate) with other atoms of the same type, (ii) consistent orientation propagation across shared boundaries, and (iii) composition in finite steps into closed simplicial structures of minimum order $K = 7$.

These are necessary conditions, not preferences. A structure failing (i) cannot participate in pairwise simplicial face-sharing. A structure failing (ii) cannot carry the oriented distinction-content required for record-supporting closure. A structure failing (iii) cannot reach $K = 7$ even with arbitrarily many composition steps. The criteria are not external filters applied to candidates; they are internal requirements of what $K = 7$ simplicial closure *is*.

Evaluating the five canonical candidate substrate realizations against (i)–(iii):

- *Half-plane*. Fails (i). A half-plane has an infinite linear boundary with no privileged sub-elements. Pairwise sharing of codimension-1 sub-elements is not well-defined without additional structure to demarcate segments of the boundary, and the half-plane does not supply that structure from its own definition.
- *Oriented disk*. Fails (iii). Closed boundaries cannot hinge pairwise in the way simplicial faces require. Two disks sharing a closed boundary coincide on the whole boundary, producing a degenerate identification rather than an edge-sharing composition into a larger complex.
- *Crease*. Fails (ii). Absent orientation, consistent orientation propagation is trivially undefined — there is no orientation to propagate across the shared boundary. A crease network can form closures but cannot carry the side-information that §3.3 requires.
- *Bordered manifold*. Fails (i) for the same reason as the half-plane, generalized. A bordered manifold has a boundary as a topological feature but does not, by its own definition, supply privileged codimension-1 sub-elements along which pairwise sharing with other bordered manifolds would be well-defined. Additional structure (a simplicial decomposition, a specified sub-boundary atlas) would have to be imposed from outside the bordered-manifold definition, at which point the resulting object is no longer just a bordered manifold but a bordered manifold *plus the additional structure* — which reduces to the fold when that structure is the minimal one required for composition.
- *Fold*. Passes all three. Folds share codimension-1 boundary segments pairwise by construction — the shared boundary is the segment along which they hinge (criterion i); they carry orientation through composition because the marked-region structure is preserved under pairwise identification of boundaries (criterion ii); and they compose into closed simplicial structures at minimum order $K = 7$ — the fold-composition relation realizes the simplicial face relation by construction of the fold definition in §4, and the minimum closure order $K = 7$ is established in the no-go companion paper (criterion iii).

The argument is therefore not "the fold happens to fit, so we pick it." It is: any substrate realization of the minimum distinction that fails (i)–(iii) cannot compose at $K = 7$, and among the canonical candidates only the fold satisfies all three. The criteria are necessity criteria, and the fold is selected by necessity among the considered candidates. §4.3 formalizes this as the *Uniqueness of the fold under dual constraints* theorem, deriving fold structure directly from four admissibility constraints (C1)–(C4) that distil the necessity conditions above. The theorem's

logical scope is any primitive satisfying the constraints; the §4.2 classification enters only as a cross-check against that theorem, not as a limitation on its scope.

4.3 Uniqueness of the fold under dual constraints

The structural arguments of §§3–4 establish that any primitive serving as the minimum actualized distinction must satisfy four conditions. This subsection formalizes those conditions and states the uniqueness result as a theorem, closing what §4.2 left as a gap between selection among canonical candidates and full uniqueness under the constraints.

Definitions

- *Distinction structure.* A triple $D = (R_1, R_2, \partial)$, where R_1 and R_2 are disjoint relata and ∂ is the locus of separability between them.
- *Minimal distinction.* A distinction structure (R_1, R_2, ∂) is minimal if neither R_1 nor R_2 contains a proper sub-region distinguished from its complement within that region. This ties directly to §3.1's binarity argument: any proper sub-distinction would itself be a binary distinction, and its presence would make the containing structure composite rather than minimal.
- *Oriented distinction.* A distinction equipped with a mapping $\omega : \{R_1, R_2\} \rightarrow \{\pm 1\}$ such that $\omega(R_1) \neq \omega(R_2)$, encoding commitment asymmetry. Equivalently, orientation can be encoded as a directional asymmetry on ∂ (the marked-region asymmetry of §3.3); the ω formulation is the informational encoding and the directional formulation the geometric one, and they are interchangeable.
- *Boundary-realized distinction.* A distinction whose relata R_1, R_2 are regions of a substrate and whose ∂ is a codimension-1 interface separating them.
- *Closure-admissible primitive.* A primitive U supporting pairwise codimension-1 boundary sharing, consistent orientation propagation under composition, and finite composition into closed simplicial structures at minimum order $K = 7$.

Admissibility constraints

- (C1) *Minimality:* U is a minimal distinction.
- (C2) *Boundary realizability:* U is boundary-realized.
- (C3) *Orientation:* U is oriented.
- (C4) *Closure admissibility:* U is closure-admissible.

These constraints are the formalization of what §§3.1–3.3 and §4.2 argue structurally. (C1) captures §3.1 (minimum distinction is binary and non-decomposable). (C2) captures §3.2 (boundary is required for separability). (C3) captures §3.3 (orientation is required for commitment). (C4) captures §4.2 (composition into $K = 7$ closure is the substrate-level admissibility condition, distilled from the companion no-go theorem).

Theorem (Uniqueness of the fold under dual constraints)

Let U be a primitive satisfying (C1)–(C4). Then U is isomorphic, up to relabelling of substrate regions, to a fold:

$$\mathbb{S}U \cong \langle R_+, R_-, \partial \rangle$$

where R_+ and R_- are substrate regions identified along a shared boundary ∂ , with orientation inherited from commitment. The isomorphism preserves boundary structure and the orientation mapping $\omega : \{R_+, R_-\} \rightarrow \{\pm 1\}$.

The theorem is a conditional: existence of such U is not separately proved, since the fold itself witnesses existence. The theorem's content is that any primitive satisfying (C1)–(C4) reduces to this witness up to substrate-region relabelling.

Proof

We proceed in four steps, each invoking one admissibility constraint to derive a structural feature of U .

Step 1 — Reduction to binary structure (from C1). By (C1), U contains no decomposition-prior distinctions. From §3.1, any non-binary distinction decomposes into binary separations plus binary consistency distinctions. Therefore U is binary: it has exactly two relata.

Step 2 — Boundary realization (from C2). By (C2), U is realized as a separation of substrate regions. U is therefore not a purely abstract relation; it takes the form of two regions with a separating codimension-1 interface. Hence $U = (R_1, R_2, \partial)$, with R_1, R_2 disjoint substrate regions and ∂ their shared boundary.

Step 3 — Orientation (from C3). By (C3), U carries directed asymmetry. R_1 and R_2 are not interchangeable: one side is structurally distinguished as committed, the other as complement. Relabelling $R_1 \rightarrow R_+$ and $R_2 \rightarrow R_-$, we have $U = (R_+, R_-, \partial)$ with intrinsic orientation $\omega : \{R_+, R_-\} \rightarrow \{\pm 1\}$.

Step 4 — Composition admissibility (from C4). Steps 1–3 have established that $U = (R_+, R_-, \partial)$ with orientation ω . (C4) adds that this structure must admit pairwise codimension-1 boundary sharing, consistent orientation propagation under composition, and finite composition into closed simplicial structures at minimum order $K = 7$ — exactly the composition behaviour specified in the fold definition of §4.

Therefore U satisfies the §4 definition of a fold in full, up to relabelling of substrate regions $R_1 \rightarrow R_+, R_2 \rightarrow R_-$. The isomorphism preserves boundary structure and the orientation mapping ω .

■

Cross-check against canonical candidate classes

The proof above is constructive: Steps 1–4 derive fold structure from (C1)–(C4) directly. As a corroborating check — not a step of the proof — one can ask which of the candidate substrate

realizations considered in §4.2 satisfies all four constraints. Purely relational primitives (graph edges, ordered pairs, abstract asymmetric relations) fail (C2). Extended region primitives (half-plane, oriented disk, bordered manifold) fail (C4): the half-plane and bordered manifold lack discrete boundary-sharing structure; the oriented disk has closed boundaries producing degenerate pairwise identifications. Non-oriented crease structures fail (C3). Only the fold class satisfies all four constraints simultaneously. The derivation above shows this directly; the cross-check corroborates it from the other direction, reading the theorem's content against the candidate classification of §4.2.

This cross-check runs against the *canonical partition* — the four structural classes identified in §4.2. Whether that partition is itself exhaustive over all admissible primitives is a separate question concerning the classification of substrate realizations, addressed in the Remarks below and flagged as an open item in §9. The proof of the theorem does not depend on partition exhaustiveness: Steps 1–4 apply to any U satisfying the constraints, regardless of which canonical class (if any) U belongs to.

Nature of the result

The theorem establishes structural uniqueness under constraints (C1)–(C4), uniqueness up to isomorphism of substrate regions, and necessity rather than stipulation. It does *not* claim uniqueness across all conceivable ontologies (the constraints encode the programme's specific commitments about minimum distinction, boundary realization, orientation, and $K = 7$ closure) and does not claim derivation independent of those constraints. What it does claim is that *within* the programme's commitments — themselves argued structurally in §§2–4 — the fold is the uniquely admissible primitive, by construction from the constraints.

A consequence worth stating explicitly: any counterexample to the theorem must exhibit a primitive U such that the constructive inferences of Steps 1–4 fail for it. This requires either (a) a primitive that violates one or more of (C1)–(C4) — in which case it lies outside the theorem's scope and is not a counterexample, or (b) a challenge to the specific structural inferences of Steps 1–4 themselves, which rest on the arguments of §§3.1, 3.2, 3.3, and §4.2. No third option is available. A separate and independent question — whether the canonical classification of §4.2 used in the cross-check is exhaustive over all substrate types — does *not* bear on the theorem's validity, only on the scope of the cross-check; it is addressed as an open item below.

Corollary (Dual-derivation consistency)

The fold is the unique primitive satisfying both the minimal-asymmetry requirement (from the regress constraint of §2) and the boundary-realization requirement (from the void constraint of the companion Fold Interface Law). That the two derivation routes converge on the same object follows necessarily from the uniqueness theorem: any primitive satisfying (C1)–(C4) is isomorphic to a fold, so any primitive produced by either derivation route — each of which entails a subset of the constraints that, combined with the other subset, yields (C1)–(C4) — must coincide with the fold produced by the other. The convergence of the two routes, argued in §4 as the intersection of two necessity conditions, is therefore a theorem rather than a coincidence.

Remarks

The theorem's logical content is that (C1)–(C4) jointly imply fold structure, by the constructive derivation of Steps 1–4. The cross-check against canonical candidate classes is a corroborating exercise that reads the theorem's content against the §4.2 classification; it is not a logical step of the proof.

A separate residual question concerns the completeness of the canonical classification itself. Whether exotic primitives outside the four canonical classes (relational / extended / non-oriented / fold) exist is a question about the typology of substrate realizations, independent of the theorem. If such exotic primitives exist and satisfy (C1)–(C4), the constructive proof applies to them — Steps 1–4 derive fold structure from the constraints regardless of which canonical class the primitive falls into, so exotic satisfiers of the constraints are still folds by the theorem. If such exotic primitives exist and *fail* one of (C1)–(C4), they lie outside the theorem's scope. The open question is whether the §4.2 classification exhausts the substrate realizations that arise in the programme, not whether the theorem applies to them.

A category-theoretic formulation — expressing the fold as the universal object in the category of primitives satisfying (C1)–(C4) — would make the theorem's scope categorically transparent and is flagged in §9 as the principal mathematical target for follow-up. The present theorem formalizes what §4.2 argued by elimination, with the cross-check preserving the connection to that earlier classification.

4.4 Pre-metric status

A further clarification. The fold is *sub-metric*: it does not have a length. Length, area, and the Planck scale are quantities that arise once many folds have accumulated into certifiable records — in particular, the Planck scale marks the floor of metric admissibility, arising at the bridge between fold dynamics and metric structure, and the closure scale ξ at which $K = 7$ structures stabilize into persistent fact-production sits at a third, higher structural level. Fold (ontological minimum), Planck scale (metric floor), and ξ (coherence threshold) are distinct and must not be conflated. The relation between them is developed in the companion paper "Fold, Planck, and Closure"; for present purposes, the point is only that the fold is not a Planck unit and has no intrinsic metric content.

5. The informational aspect: the bit

Seen informationally, the same structure is a **bit**.

A bit is the minimum information content: the resolution of a single binary alternative. It has two possible values; the actualized bit is one value rather than the other; and its values are not interchangeable — "0" and "1" are directed, not symmetric. The correspondence with the fold is exact: the two sides \leftrightarrow the two possible values; the boundary \leftrightarrow the resolution of the alternative; the orientation \leftrightarrow the directedness.

One clarification about vocabulary. The bit here is *structural* rather than ensemble-theoretic. A Shannon bit is defined relative to a prior probability distribution over an ensemble of alternatives; the bit of the minimum distinction has no such prior ensemble, because the two alternatives are constituted *by* the distinction and not prior to it. The correspondence between the structural bit and the Shannon bit runs through shared structure — binary alternative, resolution, directedness — not through a shared probabilistic setup.

6. Fold and bit as dual descriptions

The fold and the bit are not two independently motivated primitives that happen to line up. They are two descriptions of the same minimum: that structure admits one natural description in terms of its substrate feature (fold) and one in terms of its resolved content (bit), and each description specifies the same underlying object completely.

6.1 The mapping at the unit level, and where the real duality lives

A compact statement of the mapping. Let F denote a fold: an ordered pair $\langle R_+, R_- \rangle$ of substrate regions sharing a boundary ∂ , with R_+ marked as the committed side. Let $B = \{0, 1\}$ denote the bit alphabet with a fixed convention ($1 \leftrightarrow$ committed, $0 \leftrightarrow$ complement).

At the level of a single unit, the mapping $\phi : F \rightarrow B$ is a relabelling: given the convention, $\phi(\langle R_+, R_- \rangle) = 1$. This is structurally trivial — both descriptions carry exactly the same information, and the mapping does no work by itself beyond naming conventions. It is worth stating the triviality plainly: the substantive duality claim is *not* at the unit level. This is also consistent with §4.1: at the unit level, the fold and the bit descriptions (along with graph-edge, abstract-relation, and oriented-pair descriptions) are all interchangeable, because they all specify the same binary-with-orientation content.

The substantive duality claim is at the level of *composition*. Two folds F_1, F_2 sharing a sub-boundary compose to a joint substrate structure $F_1 \bowtie F_2$. Their orientations either agree or disagree along the shared sub-boundary. If they agree, the joint structure inherits a consistent orientation, and the bits combine multiplicatively under the convention. If they disagree, the joint structure has an orientation flip at the shared sub-boundary, and the bit combination carries a sign change.

The composition correspondence, as a formal target. The claim can be stated as a formal target even when full proof is deferred. Let \mathcal{F} denote the fold composition category: objects are folds, morphisms are composition operations along shared sub-boundaries. Two composition operations are defined, indexed by a sub-boundary ∂' :

- $F_1 \bowtie_{\partial'} F_2$ — *agreeing-orientation composition*, defined when the orientations of F_1 and F_2 along ∂' are consistent. The result is a joint substrate structure carrying a single consistent orientation.
- $F_1 \bowtie_{\partial'}^* F_2$ — *opposing-orientation composition*, defined when orientations oppose along ∂' . The result carries an orientation flip localized at ∂' .

Let \mathcal{B} denote the bit composition algebra on $B = \{0, 1\}$, enriched with an orientation marker to distinguish agreeing from opposing combinations. Two composition operations on \mathcal{B} correspond: $b_1 \cdot b_2$ for agreement, and $b_1 \cdot^* b_2$ for opposition (with the sign/flip carried by the marker).

The duality claim is that the mapping $\varphi : F \rightarrow B$ extends to a *structure-preserving correspondence* $\tilde{\varphi} : \mathcal{F} \rightarrow \mathcal{B}$ such that:

$$(1) \tilde{\varphi}(F_1 \bowtie \{\partial'\} F_2) = \tilde{\varphi}(F_1) \cdot \tilde{\varphi}(F_2) \quad (2) \tilde{\varphi}(F_1 \bowtie^* \{\partial'\} F_2) = \tilde{\varphi}(F_1) \cdot^* \tilde{\varphi}(F_2)$$

at every composition order, with closure preservation as a structural theorem — i.e., a composition of K folds is closed in \mathcal{F} if and only if the composition of the corresponding K bits is closed in \mathcal{B} , for every K including $K = 7$.

What this paper establishes: (i) the mapping φ is well-defined at the unit level; (ii) the composition operations on both sides are well-defined at order 2, as specified above; (iii) equations (1)–(2) hold at order 2, up to the convention fixed in §4; (iv) the correspondence extends to higher orders in a form amenable to full formalization. Moreover, the structure of the correspondence at order 2 constrains higher-order composition uniquely up to naturality conditions — i.e., any admissible extension of $\tilde{\varphi}$ to order n is determined by its action at order 2 together with compatibility with fold composition at intermediate orders, leaving no freedom in the form of the extension beyond naturality. The *full proof* — specifically the naturality conditions guaranteeing (1)–(2) at arbitrary composition order, and the structural theorem on closure preservation — is the target of a companion technical paper. The load-bearing claim of §6 is the existence of this correspondence at the level specified, not the trivial unit-level relabelling.

6.2 Implications of the duality

The composition correspondence (sketched, to be formalized) carries three implications, stated carefully:

- Any derivation carried out in the geometric vocabulary (fold composition, topological closure) is *expected*, on the composition correspondence, to have an informational counterpart (bit combination, information closure), and vice versa. Whether this can be promoted to a *general translation theorem* — guaranteeing loss-free translation at every composition order — is the open question §9 flags.
- Physical quantities that appear geometric (areas, volumes, curvatures) and those that appear informational (entropy, channel capacity, bit counts) are, on this view, expressions of the same underlying structure at the level of minimal units. Their difference at higher levels is a difference of *accumulation*, not of primitive.
- The compatibility of information-theoretic and geometric derivations of fundamental quantities — black-hole entropy being the canonical example — is *consistent with* the fold-bit duality. The duality, if it holds generally at the composition level, would explain that compatibility. This is noted as a consistency point, not as independent evidence.

The duality is load-bearing at the composition level. It is part of why independent derivation routes within the programme — some framed geometrically, some informationally — are expected to converge on compatible yields.

7. What is built from the minimum: distinction, fact, record

A single fold / bit is the smallest distinction, but it is not yet a fact. A fact requires *closure* — a structure adequate to support a stable, irreversibly committed record. The minimum structure capable of this is $K = 7$, established in earlier work as the smallest simplicial relational substrate admitting fact-production.

The relation is compositional:

- A *distinction* is a single fold / bit — the minimum asymmetry.
- A *fact* is a closed structure of distinctions — the minimum configuration ($K = 7$) in which distinctions hold together stably.
- A *record* is a fact that participates in further dynamics — a closure that persists and influences what comes next.

7.1 Two foundational levels, not in competition

The distinction between the fold/bit and the $K = 7$ closure is structurally important and easy to confuse. The regress argument of §2 establishes the fold/bit as the *fundamental unit* — the atom of actualized reality. $K = 7$ is the *minimum closure* — the atom of fact-production. These are foundational at different levels and answer different questions:

- *Fundamental unit*. What is actualized reality *made of*? The minimum distinction (fold/bit).
- *Minimum fact*. What is the smallest *stable, closed, record-supporting* structure built from the fundamental units? The $K = 7$ closure.

Conflating them generates errors in both directions. Treating $K = 7$ as the fundamental unit would fail the regress ($K = 7$ contains distinctions as parts). Treating the fold/bit as the minimum fact would fail the closure requirement (a single fold does not close). The programme's architecture requires both levels, each in its proper role. This paper's business is the lower level; the $K = 7$ closure is developed in its own companion papers.

8. Location of the unit relative to observables

The paper is pre-physical — fold and bit are sub-metric objects, and no observable quantity is extracted at this level. But the bridge to the programme's measurable content is short, and worth stating so the unit is not left orphaned from physics.

Geometric quantities arise from fold accumulation. Area, volume, and curvature are counted by the composition structure of accumulated folds, once enough folds have closed into records for

metric geometry to be admissible — specifically, once the closure scale ξ has been reached, at which $K = 7$ structures stabilize into persistent fact-production (see the "Fold, Planck, and Closure" companion paper for the derivation of ξ and its relation to the Planck floor of metric admissibility). The rich geometric structure invoked in later programme papers — \mathbb{C}^4 representations of fold composition, $\mathbb{C}\mathbb{P}^3$ projective geometry at the physical scale, curvature and internal symmetry structure — is built on this accumulation through a chain developed in the geometry papers of the programme: *minimal fold* $\rightarrow \mathbb{C}^4$ *representation of oriented composition* $\rightarrow \mathbb{C}\mathbb{P}^3$ *projective quotient at physical scale*.

The first link in this chain can be stated more precisely. The fold's binary-plus-orientation structure established in §§3–4 maps uniquely to \mathbb{C}^4 under the minimal reversible representation developed in the Fold Interface Law: the two-side \times two-orientation-phase product provides the four orthogonal basis states of the \mathbb{C}^4 representation, and the reversibility constraint fixes the representation up to a unitary of \mathbb{C}^4 . This closes the loop from the unit-level structure of §§3–4 to the $\mathbb{C}^4 \rightarrow \mathbb{C}\mathbb{P}^3$ geometric chain, and from there to the rich geometry of the later programme papers. The point here is only that this rich geometry is *downstream* of fold composition, with the first bridging step (fold $\rightarrow \mathbb{C}^4$) now explicitly anchored rather than left implicit.

Informational quantities arise from bit accumulation. Entropy is the logarithm of the accumulated bit-state count under the composition rules; in the Single-Source formulation of the programme, entropy is a functional of the committed record density $\rho(x,t)$, which is itself a functional of accumulated fold/bit composition.

The physical constants \hbar and c enter at the level of the kinematics of fold composition — c from the causal structure of oriented composition, \hbar from the granularity of the commitment rate. Both derivations are given in companion papers of the programme rather than completed here. The Planck scale enters at the level of record certification: it is the minimum scale at which $K = 7$ closure produces a stable fact in physical spacetime, and is therefore a *derived* quantity arising from fold accumulation rather than identical to the fold. (The relation between fold, Planck scale, and the closure scale ξ is developed in "Fold, Planck, and Closure.")

That the derivational programme proceeds by counting and combinatorics — fold counts, bit budgets, closure rules — is not a methodological preference. It is what is available once the fundamental unit is fixed by the regress argument. A programme built on continuous fields would derive things differently, and would, by the argument of §2, be recoverable from a more fundamental programme built on minimum distinctions. The derivational style is forced by the nature of the unit.

A concrete note on what this style has delivered. The fold/bit framework is not argued here in a vacuum — the broader programme reports a number of derivational results that rest on this foundation. The fine-structure constant is derived to $\alpha^{-1} \approx 137.034$, agreeing with measurement at roughly 15 parts per million (six-paper series). $K = 7$ is derived from six independent structural routes that converge on the same value ($K = 7$ convergence paper). The cosmological constant is derived via the Two-Planck Principle along three convergent routes. The Standard Model gauge structure $SU(3) \times SU(2) \times U(1)$, the weak mixing angle, the Higgs mass, and quark masses are derived from hexagonal geometric structure (Standard Model from Hexagonal

Geometry paper). The κ -field mass coefficient $C_m = \sqrt{(4/3)}$ arises from $PGL(3,2)$ irreducibility on the 6D non-uniform subspace V_6 . The closure scale ξ is derived as a structural consequence of three axioms rather than left as a free parameter. Full derivations are presented in the referenced programme papers listed in the References section; this paper does not reproduce them. These results are listed here not to rehearse derivations developed elsewhere but to anchor the foundational argument in a track record: the fold/bit is not a philosophical primitive in search of a physics, but the structural primitive that the programme's derivations of observable quantities presuppose.

9. Claimed, assumed, open

Claimed (uniqueness of the unit). The fundamental unit of actualized reality must be the minimum distinction. Any larger candidate contains distinctions as parts and fails the regress; below the minimum distinction is not a more fundamental distinction but the Void, which is not a unit at all. The two-option exhaustion forces uniqueness, conditional on well-foundedness of decomposition (§2.1, §2.3).

Claimed (robustness against holistic alternatives). The principal alternative ontologies that reject parts-first mereology — field-theoretic, entanglement-holistic, relations-first — still operate through distinctions at the level of their content. The regress applies to distinction-content, not to spatial mereology, and survives translation into each of these ontologies. Defeating the regress requires the much stronger claim that the alternative ontology has no decomposable content at all (§2.5). Relational-first ontologies specifically face a trilemma (content-by-relata \rightarrow regress applies; content-without-relata \rightarrow ontology vacuous; content-by-meta-relations \rightarrow second-order regress) from which no current formulation we are aware of escapes (§2.5).

Claimed (robustness against specific research programmes). The regress survives encounter with the principal live fundamental-physics programmes. Wheeler's "it from bit" is compatible (a precursor, not a rival). AQFT, loop quantum gravity, causal set theory, and constructor theory each identify a fundamental structure at their own level whose content is delivered through distinctions; the regress applies at the distinction level, placing each programme correctly above rather than below the minimum distinction. String theory is ruled out as fundamental by the extension argument. No currently active fundamental-physics programme identifies a primitive that is not itself a structure of distinctions (§2.6).

Claimed (level distinction). The fundamental unit (minimum distinction, fold/bit) and the minimum fact ($K = 7$ closure) occupy distinct structural levels and must not be conflated (§7.1).

Claimed (structurally, from distinction alone). A minimum distinction has two sides and a boundary. The binary character follows from the decomposition of n-ary partitions into binary separations plus binary consistency distinctions (§3.1). The boundary follows from the requirement that separability have a locus (§3.2).

Claimed (structurally, from commitment). A minimum *actualized* distinction is oriented. Orientation is forced by the requirement that commitment produce records, tightened by two

further links: symmetry kills discrimination (no information without asymmetry), and orientation is the structural correlate of irreversibility (§3.3).

Claimed (abstract universality at the unit level). At the unit level, the fold description is not privileged: several binary-oriented primitives — directed graph edge, asymmetric binary relation, abstract oriented pair — are abstractly isomorphic to the fold. What privileges the fold is its substrate realization, not its unit-level abstract structure (§4.1).

Claimed (necessity, composition level). The criteria (i)–(iii) required for $K = 7$ simplicial closure are necessity conditions, not preference filters: they are what "composing into $K = 7$ closure" structurally means. Among the canonical candidate substrate realizations, only the fold satisfies all three (§4.2).

Claimed (uniqueness under constraints, by theorem). Formalized as the *Uniqueness of the fold under dual constraints* theorem of §4.3: any primitive U satisfying the admissibility constraints (C1) minimality, (C2) boundary realizability, (C3) orientation, and (C4) closure admissibility is isomorphic, up to relabelling of substrate regions, to a fold — derived constructively in four steps. The theorem includes as a corollary that the two-route convergence between the regress derivation and the void-boundary derivation is necessary rather than coincidental. A cross-check against the §4.2 canonical classification corroborates the theorem against the four canonical classes. Whether the canonical classification is itself exhaustive over all substrate types is a separate open question (see §9 Open) and does not bear on the theorem's validity.

Claimed (duality). At the unit level, this duality is trivial — the fold and bit descriptions carry identical information, and the mapping between them is a relabelling under convention (§6.1). The substantive claim is at the composition level: the existence of a composition correspondence $\tilde{\varphi} : \mathcal{F} \rightarrow \mathcal{B}$ between fold composition and bit composition that preserves orientation and closure. Stated formally: $\tilde{\varphi}(F_1 \bowtie \{\partial'\} F_2) = \tilde{\varphi}(F_1) \cdot \tilde{\varphi}(F_2)$ for agreeing-orientation composition, and $\tilde{\varphi}(F_1 \bowtie^* \{\partial'\} F_2) = \tilde{\varphi}(F_1) \cdot^* \tilde{\varphi}(F_2)$ for opposing-orientation composition, with closure preservation at every order. Established at order 2 and amenable to full formalization; full proof with naturality at arbitrary order is the target of a companion technical paper (§6.1).

Claimed (pre-metric). The fold is sub-metric. The Planck scale is a derived quantity arising from fold accumulation, not identical to the fold (§4.4).

Claimed (two-route convergence via intersection). The fold is derived by two independent routes in the programme: the informational/structural route of this paper (minimum distinction \rightarrow binary + boundary + orientation \rightarrow fold) and the topological/geometric route of the Fold Interface Law (void + commitment boundary + topology \rightarrow fold). These are not two different objects but the same object, and the convergence is not accidental: the void constraint forces boundary structure, the regress forces minimal asymmetry, and the fold is the object in the intersection of both necessity conditions. The convergence is structural, not fortuitous, and parallels the multi-route derivations of α , $K = 7$, and the cosmological constant established in their respective papers (§4).

Assumed (distinctions as content of the actualized). That reality is composed of distinctions (Premise 1 of the regress). Imported from the layered foundation and not argued here. The regress argument and everything that follows are conditional on this.

Assumed (parts prior to wholes, in the decomposition-content sense). That decomposition-constituents are prior to the wholes they compose (Premise 3 of the regress). The standard mereological principle, defended at §2.5 against the principal holistic objections. A fully-worked relational-first alternative remains a genuine open question, noted below.

Assumed (well-foundedness of decomposition). That the relation "has-as-a-distinction-part" is well-founded on actualized distinctions — i.e., there is no infinite descending chain. If this fails, there is no fundamental unit at all, and the regress argument establishes only a conditional: *if* there is a fundamental unit, it must be the minimum distinction.

Assumed (irreversible commitment). That the minimum asymmetry is *irreversibly* committed rather than a fluctuation that un-commits. Imported from the actualization principle of the companion paper. Reversible minimum asymmetries would yield a different framework.

Open.

- *General translation theorem.* Whether the fold-bit composition correspondence can be promoted to a theorem showing loss-free translation between geometric and informational vocabularies at every composition order. Sketched in §6.1; full formal equivalence is a target.
- *Uniqueness beyond the canonical partition, closed via categorical formulation.* The §4.3 theorem derives fold structure from (C1)–(C4) constructively; its cross-check against the §4.2 canonical classification shows that the four canonical classes (relational / extended / non-oriented / fold) are consistent with the theorem's content. Whether the canonical classification itself is exhaustive — whether exotic substrate realizations could exist outside the four canonical classes — is a separate open question about the typology of substrate realizations. The natural closure route is a category-theoretic formulation expressing the fold as the universal object (initial or terminal, depending on framing) in the category of primitives satisfying (C1)–(C4); such a formulation would make the theorem's scope categorically transparent without relying on the canonical partition. This open item also subsumes the earlier "derivational fold at full generality" question — whether the stipulative §4 definition of the fold can be replaced entirely by a derivational identification — since Steps 1–4 already derive fold structure from the constraints, and a categorical formulation would extend that derivational identification to all admissible primitives uniformly. The alternative route — exhibition of an exotic counter-example that satisfies (C1)–(C4) but falls outside the canonical partition — is consistent with the theorem (Steps 1–4 would still apply) and would only refine the typology rather than refute the uniqueness result. This is the principal mathematical target for follow-up.
- *Full composition correspondence.* The composition rule between fold composition and bit composition is sketched in §6.1; its full formalization, with naturality conditions at arbitrary order, is deferred.

- *Escape from the relational-first trilemma.* §2.5 formulates a trilemma that any relational-first ontology must navigate (content-by-relata \rightarrow regress applies; content-without-relata \rightarrow ontology vacuous; content-by-meta-relations \rightarrow second-order regress). No current formulation we are aware of escapes all three horns, and the burden of doing so is correctly located on the relational-first side. Whether a relational-first ontology can be constructed that escapes the trilemma is a genuinely open question — unresolved by this paper, and arguably unresolvable within it — but its resolution is not required for the regress argument to go through on the position the programme actually occupies.

Conclusion

The fundamental unit of actualized reality is the minimum distinction — not by stipulation, but because no other candidate survives the regress. Any larger unit contains distinctions as parts; the parts are structurally prior; the candidate fails. The regress terminates uniquely at the minimum distinction, because below it lies the Void, which is not a more fundamental distinction but the absence of distinction. The regress applies to decomposition-content rather than to spatial mereology, and survives translation into the principal holistic alternatives — field-theoretic, entanglement-holistic, relations-first — each of which still operates through distinctions at the level of its content. Conditional on well-foundedness and on the layered-foundation commitment that actualized reality is composed of distinctions, this forces the minimum distinction as the unique candidate within any well-founded distinction-based ontology.

What that minimum is, structurally, is then fixed by three arguments. Binariness follows from the decomposition of non-binary partitions into binary separations and binary consistency conditions; boundary-structure from the requirement of separability; orientation from the requirement that commitment produce records, tightened by the symmetry-kills-discrimination and orientation-as-irreversibility links. At the unit level, the fold description is one of several interchangeable descriptions of binary-with-orientation content. What privileges the fold is its substrate realization, selected by *necessity criteria* — not preference — on composition into $K = 7$ closure. This selection is now formalized as the §4.3 uniqueness theorem: any primitive satisfying four named admissibility constraints (minimality, boundary realizability, orientation, closure admissibility) is isomorphic to a fold up to substrate relabelling, derived constructively in four steps. The theorem transforms what was previously an argument by elimination into a formally derived necessity statement under explicit admissibility constraints, with the residual open question narrowed to the completeness of the canonical classification used in the cross-check — a question targeted for closure by a category-theoretic formulation as future work. At the unit level the fold/bit descriptions coincide trivially; the substantive duality claim is the composition correspondence, sketched here and deferred to formal treatment.

Above the fold/bit sit the closures, facts, and records of the broader programme. The accumulation of fold/bits gives rise, at higher structural levels, to the observable currency of the programme — $\rho(x,t)$, entropy, the physical constants, and the derivational yields reported in the technical papers (α^{-1} to 15 ppm, $K = 7$ from six convergent routes, the cosmological constant, the Standard Model spectrum, the κ -field mass coefficient). The derivational style of the programme,

by counting and combinatorics, is not a preference but a consequence: it is what is available once the fundamental unit is forced to be the minimum. And the foundational argument is not an a priori exercise — it is the structural foundation that the programme's track record of derivations presupposes.

Characterizing the unit correctly — establishing its uniqueness, defending the argument against holistic alternatives, testing it against live research programmes, fixing its structure, locating the duality's substantive content at the composition level rather than at the unit level, and marking what is forced versus imported versus open — is what lets the combinatorics do their work.

Appendix A: Anticipated Questions and Responses

This appendix addresses referee questions anticipated against the paper's core claims, organized by topic. Answers are aligned with the specific claims, deferrals, and open items marked in the main text.

A.1 Core claim: why this is the fundamental unit

Q1. The regress depends on "parts are prior to wholes." Why accept that, especially in modern physics where fields or entanglement are often treated as primary?

The regress does not depend on spatial mereology; it depends on decomposition of *content*. Even in field-theoretic or entanglement-based ontologies, physical claims are made via distinctions — different field modes, values, measurement outcomes, correlation patterns. Those distinctions carry the ontology's theoretical content. If a structure's content is delivered through distinctions, then those distinctions are structurally prior in the sense relevant to fundamentality (§2.5). The field's spatial "wholeness" is orthogonal to the argument. Rejecting this would require an ontology with no distinguishable states and no decomposable content at all — a much stronger claim than the usual holism.

Q2. You assume well-foundedness (no infinite descent). Why should physics obey that?

Stated explicitly as an assumption (§9). However, any physically meaningful theory must admit finite specification of states; infinite descending structure would imply no minimal description, no stable information, no operationally definable facts. Well-foundedness is therefore required for a theory to support finite distinguishability and recordable facts. If it failed, the regress would establish only the conditional: *if* there is a fundamental unit, it must be the minimum distinction.

Q3. Why must the minimum unit be binary? Couldn't nature be fundamentally ternary or higher-arity?

Any n-ary distinction decomposes into binary separations plus binary consistency conditions (§3.1). A ternary {A, B, C} requires three binary separations (A from BUC, etc.) plus binary consistency binaries (non-overlap, boundary-meeting, exhaustive coverage). An n-ary

"primitive" is a bundle of binary distinctions with constraints, not a more primitive kind of structure. A genuinely non-decomposable n -ary relation would require a primitive richer than "this vs. not-this," which already presupposes distinctions to define its arity. The binary distinction is the only non-decomposable case, making it minimal.

A.2 The fold: why this structure specifically

Q4. You admit that fold, graph edge, and relation are equivalent at the unit level. Why privilege the fold?

We explicitly do not privilege the fold at the unit level (§4.1). The fold, directed graph edge, asymmetric binary relation, and oriented pair are abstractly isomorphic at unit level. Selection occurs at the composition level: only substrate-realized folds satisfy pairwise boundary sharing, consistent orientation propagation, and closure into $K = 7$ simplicial structures. The alternative abstract realizations come with different natural composition operations (vertex identification, relational product, sequencing) that do not produce $K = 7$ simplicial closure. The fold is not selected by definition but by necessity under composition constraints (§4.2).

Q5. You only tested a few candidate structures. How do you know something else doesn't satisfy your criteria?

The §4.3 Uniqueness theorem addresses this directly, and the cleaner reading is worth emphasizing. The theorem derives fold structure constructively from the admissibility constraints (C1)–(C4) in four steps — each step invoking one constraint to establish one structural feature of U . By the end of Step 4, U satisfies all structural conditions defining the fold of §4, regardless of which (if any) canonical class U belongs to.

The cross-check against the canonical classification of §4.2 corroborates the theorem from a different angle — showing that the four canonical classes (relational / extended / non-oriented / fold) are each consistent with the theorem's content — but is not itself a logical step of the proof. The residual open question concerns the typology of substrate realizations: whether exotic primitives outside the canonical classification exist. If such exotic primitives satisfy (C1)–(C4), the theorem still applies to them (Steps 1–4 derive fold structure from the constraints, without reference to the canonical partition). The canonical-partition issue is therefore about classification, not about the theorem's reach. The natural closure route is a category-theoretic formulation, flagged in §9 as a follow-up target.

Q6. Your criteria (i)–(iii) for $K = 7$ composition look engineered to select the fold. Why are they fundamental?

The criteria are not external filters chosen after the fact. They are a distillation of the admissibility axioms of the companion $K = 7$ no-go theorem, which establishes $K = 7$ as the minimum simplicial order admitting fact-production. Specifically:

- (i) *pairwise sharing of codimension-1 boundaries* — required for simplicial adjacency;
- (ii) *consistent orientation propagation* — required for distinguishable closure;

- (iii) *composition into closed simplicial structures at minimum order $K = 7$* — required for minimal fact formation.

Any primitive missing any one cannot participate in $K = 7$ closure by structural necessity. Explicit derivation of the distillation from the no-go theorem's axioms is deferred to a companion technical paper (§4.2); what matters here is that (i)–(iii) are inherited from an existing formal result, not invented for this paper's purposes.

A.3 Dual derivation: is this circular?

Q7. You derive the fold from both regress and void boundary. Isn't that redundant or circular?

No. The two routes impose independent constraints. The regress constrains internal structure ("what must a distinction minimally be?"); the void boundary constrains substrate location ("where can distinctions exist?"). Neither route presupposes the other. The fold is the unique structure satisfying both simultaneously (§4). This is convergence from independent necessity conditions, not circular derivation.

Q8. Why should these two derivations converge at all?

Because both constrain the same object from opposite directions. The regress forces minimal asymmetry as the fundamental unit; the void constraint forces boundary structure on any actualized distinction. Any admissible primitive must satisfy both simultaneously. The fold is the object in the intersection: the minimum asymmetry carrying the boundary structure the void constraint requires. The intersection pins the fold uniquely. This parallels the broader programme's methodological pattern, where α , $K = 7$, and the cosmological constant are each derived through multiple independent routes converging on the same value (§4).

A.4 \mathbb{C}^4 and quantum structure

Q9. You move from a 4-state structure to \mathbb{C}^4 . Why complex Hilbert space specifically?

The selection of \mathbb{C} over \mathbb{R} or \mathbb{H} is made in the companion Fold Interface Law paper, not in the present paper. The present paper only imports the result that the fold's binary-plus-orientation structure maps uniquely to \mathbb{C}^4 under the minimal reversible representation established there (§8). The Fold Interface Law argues the field selection by reversibility and minimality considerations (\mathbb{R} cannot carry required phase structure; \mathbb{H} is non-minimal relative to the four-state requirement; \mathbb{C} is minimally sufficient). The present paper defers to that companion rather than reproducing its argument.

Q10. Isn't the inner product structure imported from quantum mechanics?

Partially acknowledged. The programme shows that projection covariance requires an invariant overlap and that the minimal invariant is quadratic, motivating rather than assuming inner product structure. However, a full derivation of the complex inner product from void/fold axioms

alone is an open problem, stated as such. The present paper does not claim to derive quantum structure from the minimum distinction; it claims only to establish the primitive on which the programme's quantum reconstruction (in its own papers) operates.

A.5 Relation to existing frameworks

Q11. How does this differ from Wheeler's "it from bit"?

Wheeler identifies information as fundamental. This paper adds three things Wheeler does not supply: a regress-based necessity argument for why the informational primitive must be the minimum distinction; a structural characterization of the minimum (binary, boundary, orientation); and a geometric realization in the fold, composing at $K = 7$ closure. Wheeler's slogan is compatible and is best read as a precursor — the informational face of a more structured framework (§2.6).

Q12. Why is this more fundamental than QFT, LQG, or string theory?

These frameworks define structures whose content is delivered via distinctions: AQFT (operator algebras individuated by spectra and commutation), LQG (spin networks with discrete nodes, edges, labels), string theory (extended objects with internal positions and modes). Each presupposes distinguishability at its unit level, so each operates at a higher structural level than the minimum distinction. Placing them above the fold/bit is correct placement relative to the regress, not dismissal of their content at their own level (§2.6).

Q13. Your results (α , Λ , etc.) are mentioned but not derived here. Isn't that hand-waving?

This is a foundation paper, not a derivation paper. Its role is to justify the primitive used in the programme's derivations, not to reproduce the derivations themselves. The references to $\alpha^{-1} \approx 137.034$, $K = 7$ convergence, the cosmological constant, the Standard Model spectrum, and the κ -field mass coefficient are included to anchor the foundational argument in a track record — to show that the fold/bit primitive is not idle philosophy but underpins concrete, specific derivations of observable quantities. The derivations themselves are in the cited companion papers (§8). Foundation and derivation are distinct documentary acts.

A.6 Physical predictions and emergent structure

Q14. What experimental predictions distinguish your framework from alternatives?

The framework's testable content lives in its derivational outputs. Specifically: α^{-1} to 15 ppm is falsifiable if higher-precision measurements show disagreement outside the claimed accuracy window; $K = 7$ convergence is falsifiable if any observable derivation required $K \neq 7$; the Standard Model derivations predict specific mass ratios, weak mixing angle, and Higgs mass; the κ -field mass prediction $C_m = \sqrt{4/3}$ is falsifiable if the κ -field is detected with inconsistent mass. The foundational paper establishes the primitive; falsifiability of the programme lives in the derivation papers.

Q15. You formulate at the pre-metric level. How do you ensure Lorentz invariance emerges correctly when metric appears?

Lorentz invariance emerges at the Planck scale metric floor, where the causal structure of oriented fold composition gives rise to a light-cone structure. This is developed in the Proto-Time and Emergent Lorentz Invariance companion paper. The present paper establishes only that the pre-metric unit has the structural features (orientation, consistent propagation along fold hinges) that make emergent Lorentz invariance possible; the derivation itself is in the companion.

Q16. The closure scale $\xi \sim 10^{-5}$ m is a strange intermediate scale — thirty orders of magnitude above the Planck scale and far below laboratory scales. Why should such a scale be structurally fundamental?

ξ is not the fundamental scale. The fold is sub-metric; the Planck scale is the metric admissibility floor; ξ is the *closure stability threshold* — the scale at which $K = 7$ structures lock into persistent fact-production. Its specific value is derived from structural axioms in the ξ closure-scale paper. The thirty-order gap between Planck and ξ corresponds to a regime in which metric content is admissible but persistent coherent fact-worlds are not yet supported. This three-tier structure (fold / Planck / ξ) is a structural feature developed in the Fold, Planck, and Closure companion paper (§§4.4, 8).

Q17. You claim foundational status but show no gravitational physics. How does GR emerge?

Gravitational physics emerges through the coupling constant G at the bridge layer — the density-to-geometry coupling expressing how fold-commitment density influences emergent metric structure (from the Fold/Planck/Closure paper). Full derivation of Einstein's equations from accumulated fold dynamics is an open target in the broader programme. The present paper establishes the fundamental unit, not the gravitational superstructure; G is flagged as one of the three bridge constants (alongside \hbar and c) relating fold dynamics to emergent metric physics (§8).

A.7 Ontology and philosophy

Q18. Why assume reality is made of distinctions at all?

Not proved here; adopted as an operational requirement for physical theory. Without distinctions there are no measurements, no records, no laws. A theory that denied distinctions could not express its own denial, because the denial would distinguish itself from its negation. The framework is therefore conditional: *if physics is possible*, distinctions must exist. Flagged as an imported assumption (§9). The burden of rejecting it lies with the objector who claims physics can proceed without distinguishable states.

Q19. Does your framework require irreversibility at the fundamental level?

A distinction must be drawn. At the level of *abstract* minimum distinction (two sides plus boundary), irreversibility is not required. At the level of the minimum *actualized* distinction (the committed fold), orientation is required (§3.3), and orientation is the structural correlate of irreversibility — reversible operations leave before/after symmetric, so orientation is what commits one possibility as actual. Irreversibility is imported from the commitment principle of the companion layered-foundation paper, but once imported is tightly linked to orientation at the minimum actualized unit.

Q20. A distinction is a relation. If distinctions are primitive, relations are primitive — aren't you just doing a kind of relational ontology and pretending you're not?

Distinctions are indeed relations with specific internal structure (two relata, boundary, orientation). The difference from problematic relational-first ontologies is that the relata are *co-primitive* with the relation. A distinction is not a relation between already-individuated things (circular — relata's distinctness presupposed) nor a relation that constitutes its relata from nothing (horn b of §2.5's trilemma — relation underdetermined). It is a structure in which relation and relata are given together, neither prior. This places the programme on horn (a) of the trilemma: content is specified by what is related, and the content is well-formed because relation and relata co-emerge. The programme is a distinction-first ontology, where "distinction" names the co-emergence of a relation with its relata.

A.8 Weak points

Q21. Your composition-level duality is not fully formalized. Isn't that central?

Yes — acknowledged and marked precisely (§6.1, §9). Established: mapping ϕ well-defined at unit level; composition operations on both sides (\mathcal{F} and \mathcal{B}) defined; correspondence at order 2 specified by two numbered equations:

- $\tilde{\phi}(F_1 \bowtie_{\partial'} F_2) = \tilde{\phi}(F_1) \cdot \tilde{\phi}(F_2)$ for agreeing orientation,
- $\phi(F_1 \bowtie_{\partial'} F_2) = \phi(F_1) \cdot \phi(F_2)$ for opposing orientation;

extension to higher orders amenable to full formalization. Deferred: full proof at arbitrary composition order with naturality conditions, and structural theorem on closure preservation. Flagged as target of a companion technical paper. The load-bearing claim of §6 — existence of the correspondence — is established; full formal equivalence is the deferred result.

Q22. You claim uniqueness but admit possible alternatives. Isn't that inconsistent?

No. The paper distinguishes three levels of claim:

- *Necessity of criteria* (established, §4.2). Any primitive failing criteria (i)–(iii) cannot compose at $K = 7$.
- *Uniqueness under admissibility constraints* (established as theorem, §4.3). Any primitive satisfying constraints (C1) minimality, (C2) boundary realizability, (C3) orientation, and (C4) closure admissibility is isomorphic to a fold up to substrate relabelling. The proof is

constructive: Steps 1–4 derive fold structure from the constraints directly, independent of which canonical class the primitive belongs to.

- *Completeness of the canonical classification* (partially open). The cross-check in §4.3 runs against the four canonical classes identified in §4.2. Whether exotic primitives exist outside this classification is a separate question about substrate typology, not about the theorem. If exotic primitives satisfy (C1)–(C4) they are still folds by the theorem; the open question is only whether the canonical classification itself exhausts the types.

The current result is uniqueness *under the constraint set, by theorem*, with the remaining open question being whether the canonical classification is exhaustive. This is precise demarcation, not inconsistency. A category-theoretic formulation is the natural route to closing the residual question.

References

Programme companion papers

All items below are VERSF-programme papers by the present author, available at versf-eos.com / AIDA Institute. Titles are working titles; citation details are internal to the programme.

- Taylor, K. *The Layered Foundation of VERSF*. Imported premises: actualized reality composed of distinctions; potentiality and actualization as co-primitive aspects.
- Taylor, K. *The Fold Interface Law*. Topological derivation of the fold from the void-boundary constraint; minimal reversible \mathbb{C}^4 representation; selection of \mathbb{C} over \mathbb{R} and \mathbb{H} .
- Taylor, K. *No-Go Theorem for Non-Simplicial Relational Substrates: Derivation of $K = 7$* . Admissibility axioms; establishment of $K = 7$ as the minimum simplicial order admitting fact-production; source of the composition criteria distilled as (i)–(iii) in §4.2.
- Taylor, K. *Fold, Planck, and Closure: Structural, Metric, and Coherence Thresholds in VERSF*. Three-scale structure (fold / Planck / ξ); bridge constants c, \hbar, G ; certification functional.
- Taylor, K. *The Commitment Barrier and Closure Scale ξ* . Derivation of ξ as structural consequence of three axioms; three-axiom basis for closure stability.
- Taylor, K. *Single-Source Theorem: All Observables as Functionals of $\rho(x,t)$* . Committed record density as the observable currency of the programme.
- Taylor, K. *The Fine-Structure Constant from VERSF* (six-paper series). $\alpha^{-1} \approx 137.034$ at second order, agreeing with measurement to ≈ 15 parts per million.
- Taylor, K. *The Cosmological Constant via the Two-Planck Principle*. A derivation along three convergent routes.
- Taylor, K. *Standard Model from Hexagonal Geometry*. Derivation of $SU(3) \times SU(2) \times U(1)$, the weak mixing angle, the Higgs mass, and quark masses from hexagonal substrate structure.
- Taylor, K. *The κ -Field Mass from $PGL(3,2)$ Irreducibility*. Derivation of $C_m = \sqrt[4]{(4/3)}$ via irreducibility on the 6D non-uniform subspace V_6 .

- Taylor, K. *Proto-Time and Emergent Lorentz Invariance*. Emergence of Lorentz structure at the Planck floor from oriented fold-composition causal structure.
- Taylor, K. *Quantum Reconstruction from Fold Admissibility*. Derivation of complex Hilbert space, Born rule, Schrödinger and Dirac equations from admissibility axioms.
- Taylor, K. *Thermodynamics from Irreversible Commitment*. Connection between orientation at the minimum unit and macroscopic entropy increase.
- Taylor, K. *$K = 7$ Convergence: Six Independent Derivational Routes*. Six independent routes to $K = 7$ converging on the same value.
- Taylor, K. *VERSF Uniqueness Programme*. Master theorem on minimal fact-supporting structure.

External references

- Wheeler, J. A. (1990). "Information, physics, quantum: The search for links." In *Complexity, Entropy, and the Physics of Information*, ed. W. H. Zurek. Addison-Wesley.
- Haag, R. (1996). *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed. Springer. (Algebraic QFT.)
- Rovelli, C. (2004). *Quantum Gravity*. Cambridge University Press. (Loop quantum gravity.)
- Dowker, F. (2005). "Causal sets and the deep structure of spacetime." In *100 Years of Relativity: Space-time Structure*, ed. A. Ashtekar. World Scientific.
- Polchinski, J. (1998). *String Theory*, vols. I–II. Cambridge University Press.
- Deutsch, D. (2013). "Constructor theory." *Synthese* 190 (18): 4331–4359.
- Marletto, C. (2015). "Constructor theory of information." *Proceedings of the Royal Society A* 471: 20150883.