

Fold, Fact, Planck, and Closure: Structural, Realization, Metric, and Coherence Thresholds in VERSF

Keith Taylor — versf-eos.com

Companion to "The Minimal Asymmetry," "Why the Fundamental Unit Must Be the Minimum Distinction," "The Topological Threshold for Fact Formation," "The Planck Time as a Minimum Bit-Certification Scale," and "Multiple Structural Derivations of the Mesoscopic Coherence Scale ξ_{meso} ."

Abstract for the General Reader

What does it take for anything to be real? Before a thing can be measured, it has to first exist as an actual event rather than a mere possibility. Once it exists, its measurement has to be meaningful — there has to be a smallest scale at which "length" and "duration" make sense. And before events cohere into a stable physical world, the measurable things have to be able to hold themselves together. Each of these is a different transition, and physics tends to blur them together.

This paper sets out the four distinct levels that sit between the deepest foundations of the VERSF programme and everyday physics:

- **The fold** — the minimum possible distinction; the starting point from which everything else is built. It exists as a primitive of structure, but is not yet realized as a committed fact.
- **The minimal fact** — the first irreversible event. This is what occurs when a fold interaction becomes topologically trapped in a way that cannot be undone. It is real — something has definitely happened — but it does not yet have a well-defined size.
- **The Planck scale** — the smallest meaningful size for a realized fact. This is where length, time, and mass first become operationally admissible. About 10^{-35} metres.
- **The closure scale ξ** — the smallest scale at which facts form self-sustaining networks, rather than relying on larger surroundings to stabilize them. About 10^{-5} metres, roughly the width of a human cell.

Each level answers a different question. *What is the minimum distinction?* — the fold. *What is the first realized structure?* — the minimal fact. *When does a realized fact first acquire metric content?* — the Planck scale. *When do realized facts first form self-sustaining networks?* — the closure scale ξ . Conflating these four questions has historically caused confusion in foundational physics. This paper's main contribution is to show that the four levels are forced by the programme's other commitments — collapsing any two of them produces a contradiction — and

that the two measurable thresholds (Planck and ξ) are formally distinct rather than two descriptions of the same thing.

The paper integrates three separate derivations (of fact formation, Planck certification, and coherence closure, each worked out in companion papers) into a single hierarchy, shows why the hierarchy cannot be collapsed, and identifies the narrower open question that remains: deriving the bridge constants c , \hbar , and G themselves from the underlying substrate rather than importing them from effective physics.

Table of Contents

1. The question
2. The architecture: four levels
 - Why the hierarchy cannot be collapsed
3. The fold is not Planckian
4. The minimal fact layer: first realized structure
 - What a fact is
 - Why this is a distinct layer
 - What the minimal fact does not yet have
 - What the minimal fact does not yet do
5. How Planck units arise: the bridge principle
 - The status of c , \hbar , and G
6. The closure scale ξ : coherence, not metricity
 - Why sub- ξ physics exists: environment-mediated closure
 - A concrete case: cloud-chamber ionization
7. The certification functional: integration with the Planck Time derivation
 - Certification, defined
 - The functional
 - What the integration gives
 - Interface with quantum mechanics
8. The coherence functional: integration with the Mesoscopic derivation
 - The functional
 - Two functionals at two levels
 - Interface with general relativity
9. Planck \neq ξ : a non-equivalence proposition
 - Proposition (Non-equivalence)
 - Argument
 - Consequence
10. Claimed, assumed, open
11. Conclusion

References

Abstract

Within VERSF, the fold-bit is the fundamental unit of actualized reality: a pre-metric structural primitive, prior to length, duration, and mass. This paper sets out the hierarchical architecture that connects this primitive to physical metric content, and integrates three independent derivations into that architecture. The architecture has four structural levels. The fold is the **ontological minimum** — the least actualized distinction. The minimal fact is the **first realized structure** — the first irreversible bit, produced when fold interactions satisfy the topological trapping condition $\beta_1 \geq 1$. The Planck scale is the **metric minimum** — the floor at which an already-trapped bit first acquires certifiable length, duration, and mass, fixed by the joint constraint of quantum distinguishability and gravitational admissibility. The closure scale ξ is the **coherence minimum** — the threshold at which committed distinctions form self-supporting networks of persistent facts, fixed by competing UV/IR closure failure modes. Four levels answer four different questions: *what is the minimum distinction?*, *what is the first realized structure?*, *when does a realized fact first acquire metric content?*, and *when do realized facts first form self-sustaining networks?* Keeping the four formally distinct is required both for internal consistency and for the correct attribution of c , \hbar , G to the bridge between structure and metric. A non-equivalence proposition establishes that the Planck scale and the closure scale cannot be collapsed into a single threshold: they are fixed by different functionals over different variables, and they characterize different structural questions. The paper's distinctive contribution is to locate the three companion derivations at their proper structural levels, prevent the conflation that would otherwise damage the programme, and identify the narrower open object that remains: a substrate-level derivation of the bridge constraints that bypasses the effective causal regime entirely.

1. The question

The programme has established two things so far. The fundamental unit of actualized reality is the minimum distinction, manifest as **fold** (geometrically) and **bit** (informationally). This unit is forced by regress: any larger candidate contains distinctions as parts and is therefore not fundamental.

Throughout this paper, "fold" denotes the geometric aspect and "bit" the informational aspect of the same primitive distinction; "fold-bit" is used when both aspects are simultaneously in play. After §2 the term "fold" is used as shorthand.

The natural next question is physical rather than structural: **how does the fold relate to the Planck scale?** This is the first point at which the foundational architecture must make contact with measurable physics, and it is a point where the programme must be careful. The Planck scale is the smallest physically meaningful length, time, and mass combination, built from c , G , \hbar . It is where physics has a floor. The fold is the smallest distinction the world can support. It is where actualization has a floor. These two floors are not the same thing, and conflating them would damage the programme.

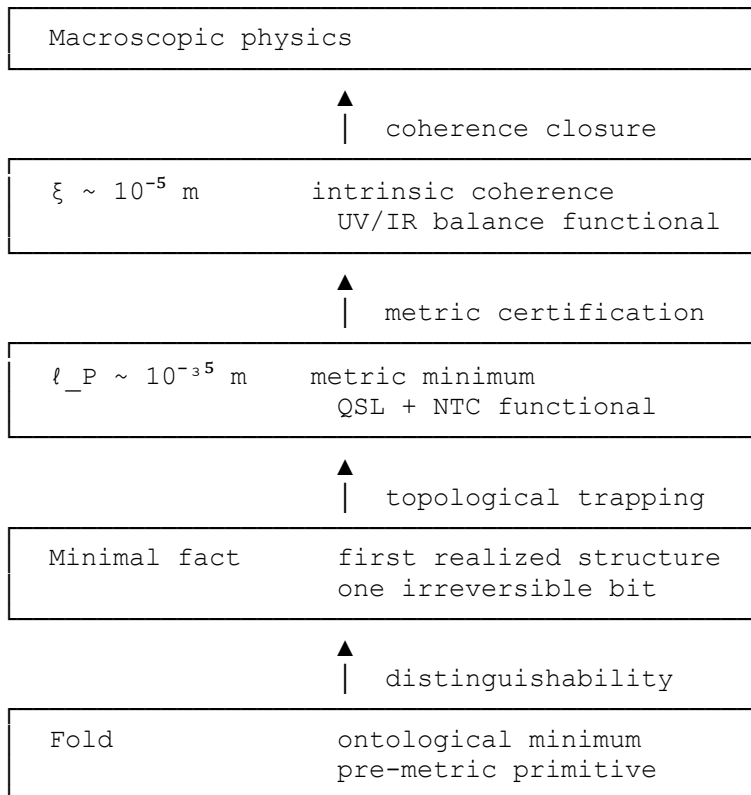
A second question follows immediately. The closure scale $\xi \sim 30\text{--}100 \mu\text{m}$ derived elsewhere in the programme sits roughly thirty orders of magnitude above the Planck scale. What is ξ a threshold *of*, if not of metric admissibility?

A third question sits between these. Between the pre-metric fold and the metrically certified Planck regime, what is the first structure that actually exists — the first thing one could point to and call *realized*?

This paper states the relationships directly. The four-question framing introduced in the abstract will recur as the structural spine.

2. The architecture: four levels

The programme has four structural levels between the pre-metric void and macroscopic physics. The first two are structural (not metric); the second two are metric thresholds. All four are distinct, and each answers a different question.



Each level answers a different question:

- The **fold** answers: *what is the minimum distinction?*
- The **minimal fact** answers: *what is the first realized structure?*

- The **Planck scale** answers: *when does a realized fact first acquire metric content?*
- The **closure scale** ξ answers: *when do realized facts first form self-sustaining networks?*

The distinction between levels 1–2 and levels 3–4 is worth naming explicitly. Levels 1 and 2 are **structural**: they concern what exists at all and what its minimum realized form looks like, without reference to length, duration, or mass. Levels 3 and 4 are **metric**: they concern the scales at which metric content first becomes certifiable (Planck) and at which coherent fact-networks first become self-sustaining (ξ). The architecture therefore separates two kinds of question that are easy to conflate: *what is there?* and *how big is it?*

The hierarchy is cumulative: the fold supplies distinguishability, the minimal fact supplies irreversibility, the Planck scale supplies metric admissibility, and the closure scale supplies persistence. Each level adds a property not present at the level below, and no level can be reached by augmenting the level below it with the wrong kind of addition. This is what gives the architecture its structural integrity — the levels are ordered not by convention but by what each one supplies that the previous one lacked.

Why the hierarchy cannot be collapsed

The four-level architecture is not a stylistic choice. Any attempt to collapse adjacent levels produces a contradiction with claims established elsewhere in the programme:

- **Collapsing fold and minimal fact (levels 1 \leftrightarrow 2)** would equate a pre-metric ontological primitive with a realized irreversible distinction. But the fold is, by construction, provisional and reversible until trapping occurs; a fact is, by construction, irreversible. Identifying them erases the distinction between what is possible and what has occurred.
- **Collapsing minimal fact and Planck (levels 2 \leftrightarrow 3)** would equate a topologically trapped bit with a metrically certified bit, breaking the pre-metric nature of the minimal fact. The Topological Threshold derivation establishes fact formation without reference to length, duration, or mass; forcing it to occur at the Planck scale imports metric content into a structural level that does not require it.
- **Collapsing fold and Planck (levels 1 \leftrightarrow 3)** would make the fold Planck-sized, which breaks the programme's core commitment that geometry is derivable from pre-metric structure rather than assumed at the primitive (§3).
- **Collapsing Planck and closure (levels 3 \leftrightarrow 4)** is ruled out by the non-equivalence proposition of §9: the two thresholds are fixed by distinct functionals over distinct input sets and acting on distinct structural objects. The UV/IR balance that fixes ξ cannot be reduced to the QSL+NTC minimax that fixes ℓ_P .

Each collapse fails at a different level, for a different structural reason. Taken together, the four failures mean the hierarchy is forced rather than chosen: the programme's other claims — pre-metric primitive, topological trapping, certification bridge, UV/IR balance — jointly require exactly these four levels and no fewer.

Each of the following sections treats one level of the hierarchy and, where applicable, the functional that fixes its associated threshold.

3. The fold is not Planckian

The fold must not be identified with a Planck unit. This is a structural claim, not a cautionary one.

Identifying the fold with ℓ_P builds metric content into the primitive from the start. **If the fold has a length in metres, geometry is already present at the foundational level, and the architecture collapses into a dressed-up Planck ontology.** That cuts directly against the programme's core commitment: geometry is downstream of distinguishability and commitment, not prior to them. The derivational fertility of the programme depends on geometry being derivable from sub-metric structure, not assumed at the primitive.

The correct statement is: **a single fold is structurally necessary but not realized.** A fold is a minimal distinction. It is not yet a realized event — that is the work of the next level — and it is not yet metrically certified — that is the work of the level after that. Planck quantities do not characterize isolated folds; they characterize the minimum metrically admissible regime in which a realized fact can be certified. The Planck scale is therefore a property of the bridge between realized facts and metric content, not of individual folds.

This distinction matters. Overloading the primitive with metric content produces a programme that is less elegant and less explanatory. Keeping the primitive clean — pre-metric, structural, dimensionless, and not yet realized — places the burden of everything that follows on the combinatorics and dynamics, which is where it belongs. The programme delivers sub-metric-to-metric derivation in two concrete places: the Topological Threshold derivation produces realized facts from fold-interaction structure with no metric inputs, and the Planck Time derivation produces metric thresholds as outputs of the certification functional rather than as primitive scales. Both would be unavailable if the fold were metric from the start.

4. The minimal fact layer: first realized structure

Between the pre-metric fold and the metric Planck regime sits a level that is structurally indispensable and often overlooked: the minimal fact. This is the first *realized* structure in the hierarchy — the first thing that has actually occurred as an irreversible event, as distinct from what is structurally present at the primitive level but not yet event-committed.

Throughout this paper "minimal fact" is the primary term for the object introduced at this level. Phrases such as "trapped bit," "first irreversible distinction," and "first realized structure" are used only as definitional elaborations of the same object.

What a fact is

A fact, in this framework, is a topologically trapped distinction. The Topological Threshold derivation establishes the condition under which trapping becomes possible: the fold-interaction graph must support cyclic connectivity, $\beta_1 \geq 1$. Tree-structured interactions cannot trap alternatives locally, because in a tree every attempt to separate two alternatives requires modifying the unique path connecting them — a global operation. Cycles provide the alternate routing that allows one alternative to enter a loop and become quarantined while the other continues; once this routing occurs, the two alternatives can never recombine.

At that moment, one bit is fixed. A distinction that was previously reversible and provisional becomes irreversible and realized. This is the minimal fact.

Stated operationally: **a minimal fact is a non-recombinable distinction that persists under reversible microdynamics and is therefore invariant under forward evolution.** The three clauses are tightly linked: non-recombinability is a topological condition on the information graph (cyclic connectivity); persistence under reversible microdynamics is what turns the topological condition into a dynamical one (trajectories do not escape the trap); invariance under forward evolution is the operational consequence (the bit stays fixed). The operational definition ties topology, dynamics, and persistence together in a single clause, and it is this definition — not any particular scale or metric — that the later levels of the hierarchy act upon.

Why this is a distinct layer

It is tempting to treat $\beta_1 \geq 1$ as merely a condition on fold interactions — a precondition for the Planck scale to apply — without recognizing that the condition *produces* something. That reading understates what is happening. $\beta_1 \geq 1$ is the enabling condition; the minimal fact is what the condition makes possible. The condition and the object produced under the condition are not the same thing: the first is a global topological property of the fold-interaction graph, the second is a local realized distinction that becomes available once the graph satisfies the property.

Concretely: the topological condition enables non-recombinability; non-recombinability produces the first irreversible distinction; that distinction is the minimal fact. Without naming the minimal fact as a distinct layer, the hierarchy has the wrong shape — it reads as *conditions* → *scales*, whereas the correct reading is *primitive* → *fact* → *metric* → *coherence*. The minimal fact is not a condition; it is the first realized object.

The stronger version of this claim is that $\beta_1 \geq 1$ does not merely permit fact formation but generically forces it. Given reversible microdynamics and local construction, any system that admits $\beta_1 \geq 1$ necessarily evolves into configurations in which non-recombinability occurs, and such configurations are dynamically unavoidable under generic evolution on any graph with $\beta_1 \geq 1$, in the sense that the measure of trajectories that avoid trapping configurations is zero under any ergodic or mixing dynamics on the graph: the cyclic connectivity provides alternate routing, local construction uses that routing to trap one alternative, and the reversible microdynamics carries trajectories into the trap whenever they encounter the appropriate configuration. The claim is not merely that trapping configurations exist in the state space — it is that the dynamics lead into them with probability one, and that avoidance is a fine-tuned exception of measure zero rather than a generic possibility. Minimal facts are therefore not merely possible at $\beta_1 \geq 1$ but

generically produced under evolution. This upgrades the hierarchy's second level from a permissive condition to a productive one, and explains why the level deserves naming: something is reliably being made here, not merely allowed.

What the minimal fact does not yet have

The minimal fact is realized but not metrically certified. At this level there is an irreversible distinction — something has definitely happened — but there is not yet a well-defined length over which it happened, a duration during which it occurred, or a mass associated with it. These are supplied by the next level. The fact is real; its metric content is not yet admissible.

To put this as directly as possible: the minimal fact is the first irreversible distinction, but it is not yet metrically certified; the Planck scale marks the first regime in which such a fact acquires operationally meaningful metric content. A minimal fact and a Planck-scale certified fact are therefore not the same object. The former is a realized distinction without a well-defined size; the latter is a realized distinction equipped with the smallest size at which "size" means anything. **The Planck scale is therefore not the scale at which facts first exist, but the scale at which facts first become metrically interrogable — the smallest scale at which questions of size, duration, and action admit invariant answers.** The distinction is clean and worth preserving: existence is the work of level 2, metric invariance is the work of level 3. Confusing the two re-imports the very conflation that §3 argued against.

This explains why the Planck scale is the right next step rather than a stipulation. Once a fact exists, the question "*how big is it?*" becomes meaningful, and the minimax between quantum distinguishability and gravitational admissibility (§7) fixes the answer. Without a realized fact to measure, the Planck scale would have nothing to be the floor of.

What the minimal fact does not yet do

The minimal fact is a single realized bit. It does not yet form a self-sustaining network. A single trapped distinction can exist without forming part of a persistent, coherent, fact-bearing structure — and in fact most single trappings at sub- ξ scales do not, which is why environment-mediated closure (§6) is needed to account for persistent micro-physics. The transition from isolated realized facts to self-sustaining networks is the work of the closure scale ξ , not of the minimal fact layer itself.

The minimal fact layer therefore sits cleanly between the fold and the Planck scale: it supplies what the fold lacks (realization) and lacks what the Planck scale supplies (metric content) and what ξ supplies (networked coherence). The layering is self-announcing for the rest of the paper: the minimal fact is the object on which the certification functional of §7 acts to supply metric content, and on which the coherence functional of §8 acts to supply networked coherence. Without the minimal fact layer, both subsequent functionals would lack a well-defined object to act on.

One further observation connects this level to the programme's broader architecture. The minimal fact is also the first unit of entropy ($k_B \ln 2$, per the Landauer derivation in the

Topological Threshold paper) and the first increment of accumulated irreversibility from which the arrow of time is constructed. The minimal fact layer is therefore not only the hinge between structure and metric in the hierarchy laid out here; it is also the origin of entropy and time in the programme's treatment of those concepts elsewhere. The four-level architecture and the entropy-and-time account are two views of the same object at level 2.

5. How Planck units arise: the bridge principle

The Planck quantities arise at the **bridge** between realized fact and physical metric. They are the first metrically meaningful composites of the dynamics that certify a realized fact.

The fold has a single structural character: distinction and commitment. When fold dynamics produce a minimal fact (§4) and that fact is then organized into a minimal closed certification process, that single character manifests physically through three distinct aspects, each associated with a fundamental constant:

- **Temporal aspect** — **c**. The limiting speed of commitment propagation. A commitment cannot influence structure faster than the rate at which commitment dynamics themselves propagate. The speed of light is the physical expression of this limit.
- **Quantization aspect** — **ħ**. The minimum action associated with irreversible commitment. Each commitment event carries a discrete, irreducible cost; that cost, expressed as action, is \hbar . The quantum of action is the physical expression of the discreteness of commitment.
- **Relational/geometric aspect** — **G**. The coupling between committed density and effective geometric response. Regions with high fold-commitment density influence effective geometry; the strength of that influence is G. Newton's constant is the physical expression of the density-to-geometry coupling.

Given these three aspects, the Planck quantities follow in the standard way:

$$\ell_P = \sqrt{(\hbar G/c^3)}, t_P = \sqrt{(\hbar G/c^5)}, m_P = \sqrt{(\hbar c/G)}.$$

The dimensional combinations are forced: the bridge involves exactly these three constants and no others, so any metric scale built from the bridge must be of the form $\hbar^p \cdot G^q \cdot c^r$ with rational exponents fixed by dimensional requirements; the three length-, time-, and mass-dimensional composites are unique up to numerical prefactor, giving $\ell_P = \sqrt{(\hbar G/c^3)}$ and its companions. The resulting statement is programme-internal but sharp. **Within the effective bridge description, the Planck scale emerges as the minimal metric output of the certification functional**, fixed by the structure of minimal closed certification processes and the three constants that characterize their physical interpretation. The effective-level qualifier matters: QSL and NTC apply as emergent constraints, so the Planck scale is derived within the effective description but still imported in the sense that those constraints are themselves imported. A fully substrate-level derivation — one that bypasses the effective regime — remains the primary open object of the programme (§10).

The status of c , \hbar , and G

The treatment above imports c , \hbar , and G from the effective physical description rather than deriving them from fold dynamics. This is a real import, and it deserves explicit acknowledgement rather than footnote treatment. The present paper does not close this gap; it locates it precisely.

The status of the three constants breaks into two parts:

Current status — imported into the effective bridge. The assignments (c as propagation limit, \hbar as action quantum, G as density-to-geometry coupling) are established elsewhere in the programme through derivations of commitment propagation, action quantization, and density-geometry coupling respectively. Those derivations themselves use effective-level structure — fold dynamics organized into processes that already have temporal ordering, discrete commitment costs, and density-response behaviour. The constants are therefore imported into the bridge at the effective level, not derived from the fold primitive directly.

Programme direction — derivable from closure relations. The programme is moving toward derivations of the constants themselves from closure conditions. The companion to the Mesoscopic Coherence paper establishes one such route: the speed of light follows from a closure relation linking the cosmological and mesoscopic scales:

$$c = (L_{IR}^2 \cdot \hbar G / \xi^4)^{1/3}$$

This expresses c in terms of \hbar , G , and scales derived elsewhere in the programme — but it does not yet derive \hbar or G . A complete closure would derive all three constants from substrate-level fold dynamics, making the Planck composites genuinely downstream of the primitive rather than constructed from imported pieces.

This is not a weakness of the present architecture; it is the architecture's principal open direction. The bridge is the right place for the constants to sit, and the architecture correctly locates them there; what remains is to derive them at the substrate level rather than importing them from the effective level. This direction is taken up explicitly in the open-questions section below.

6. The closure scale ξ : coherence, not metricity

The closure scale ξ is a separate threshold with a separate role. Confusing it with the Planck scale is a structural error the programme must avoid — and indeed the non-equivalence is formal, not merely heuristic, as established in §9 below.

The Planck scale is the floor of metric admissibility for an already-realized fact. Below ℓ_P , asking "*how long is this?*" of a realized fact ceases to have operational content. The closure scale ξ is not about metric admissibility at all; it is about **coherence stability of networks of facts**. Above ξ , commitment structure holds together intrinsically — chains of stable, interlocking

commitments that can bear records, sustain dynamics, and support the ordinary physical world without external scaffolding. Below ξ , such network closure does not stabilize *on its own*.

These answer different questions:

- **When does a realized fact first acquire metric content?** — Planck scale.
- **When do realized facts first form self-sustaining networks?** — closure scale ξ .

Derived from independent inputs, the two scales differ by roughly thirty orders of magnitude: $\ell_P \sim 10^{-35}$ m (from \hbar , G , c via the certification functional of §7) and $\xi \sim 10^{-5}$ m (from ℓ^* and L_{IR} via the coherence functional of §8). The separation is the numerical signature of their structural distinctness, not a parameter set by hand.

Why sub- ξ physics exists: environment-mediated closure

The obvious objection to the coherence-minimum reading is that physics evidently produces facts at scales far below 10^{-5} m — atomic spectra, particle tracks, single-molecule events. If these are facts, why is ξ the coherence minimum rather than an arbitrary upper cutoff?

The answer follows directly from the Topological Threshold derivation and the minimal fact layer of §4. The topological condition for fact formation ($\beta_1 \geq 1$) is a condition on the *effective information graph*, which includes the environment. Local trapping is possible at sub- ξ scales, but it is only possible because the local trap is embedded in an ambient substrate whose larger-scale connectivity provides the loop structure and the record-stabilizing machinery. The bit is fixed locally; the connectivity that makes the fixation possible is not confined to the local region.

The underlying reason is graph-theoretic, not merely physical. Because non-recombinability is defined relative to the full reachability structure of the information graph, local trapping events at sub- ξ scales depend on the global connectivity of that graph. The reachability structure is what determines whether two alternatives are in fact non-recombinable — and reachability is a global property, not a local one. A sub- ξ region can contain a local trap, but whether that trap succeeds in rendering alternatives non-recombinable depends on what the information graph looks like outside the region. Intrinsic closure is the condition under which the necessary global connectivity is guaranteed without external support; only above ξ is this condition met.

Two modes of closure must therefore be distinguished:

- **Intrinsic closure.** A commitment structure at scale L is intrinsically closed if it sustains itself as a persistent fact-bearing object without requiring support from a larger ambient structure. Intrinsic closure first becomes available at $L \geq \xi$.
- **Environment-mediated closure.** A commitment structure at scale $L < \xi$ can participate in persistent fact-production *when embedded in* a ξ -scale or larger environment that supplies the missing closure stability.

The logical chain is:

1. Local bit fixation requires $\beta_1 \geq 1$ in the information graph, including the environment (Topological Threshold).
2. Non-recombinability is defined relative to the *full* reachability structure of that graph, not the local neighbourhood.
3. Therefore local bit fixation at scale L depends on the global connectivity of the information graph at scales larger than L .
4. Therefore local bit fixation does not imply self-sustaining closure at scale L .
5. Therefore a distinct threshold — ξ — is required to characterize the scale at which closure becomes self-sustaining without external support.

Sub- ξ facts are real, but they ride on the coherent ξ -scale and macroscopic structure that envelops them. Isolated facts can occur at any scale above the Planck floor; self-sustaining fact networks require ξ . This is what allows the programme to be consistent both with the existence of microscopic measurements and with a mesoscopic intrinsic-closure threshold. The thirty orders of magnitude between ℓ_P and ξ are not a fact-free regime — they are the regime in which facts exist only through their participation in the coherent network that ξ characterizes.

A concrete case: cloud-chamber ionization

To make the environment-mediated reading concrete, consider a charged particle traversing a supersaturated vapour. The ionization event at a single vapour molecule is a sub- ξ fact: a local, irreversible distinction (ionized vs. not) occurring at atomic scale. Viewed in isolation, this event has no stable closure — the molecular state would relax back into the vapour without producing any persistent record.

What turns the ionization into a durable fact is the surrounding supersaturated vapour, which is itself a structure whose characteristic coherence scale lies at or above ξ : the droplet nucleation that follows the ionization, the $\sim 10\text{--}100\ \mu\text{m}$ droplet that grows around it, and the macroscopic vapour phase that supplies the thermodynamic instability are all ξ -scale or larger. The ionization produces a minimal fact locally, but the fact persists, propagates, and becomes checkable only because the ξ -scale environment supplies the alternate-routing global connectivity that the sub- ξ information neighbourhood alone cannot provide. Remove the supersaturation and the same ionization produces no track; the minimal fact still forms in the sense of a locally trapped distinction, but without environment-mediated closure it has no path into the record.

This is the structural shape the present paper's argument predicts: the minimal-fact condition is satisfied locally, metric certification is satisfied at the Planck floor and above, but *persistence as a fact-bearing object* — the property that lets the ionization count as an event in ordinary physics — requires the ξ -scale surroundings. The same structural pattern applies to CCD photoelectron \rightarrow pixel commitment, single-atom spectroscopy \rightarrow macroscopic spectrometer, and decay events \rightarrow detector tracks. In every case, the "microscopic fact" is really a coupled micro/meso object whose existence as a persistent record depends on the ξ -scale environment it is embedded in.

7. The certification functional: integration with the Planck Time derivation

The architecture requires a sharp mathematical condition fixing the Planck scale as the metric floor of realized facts. That condition is supplied by the Planck Time companion paper; this section integrates it.

Certification, defined

A fold process F that has produced a minimal fact is **metrically certified** when the closure of F supports a well-defined invariant length, duration, and mass under the bridge constants (c, \hbar, G) — that is, when F closes in a way that licenses operationally meaningful answers to the questions "*how long does it last?*", "*how far does it extend?*", and "*how much action does it carry?*" Certification is a property of the closure structure of F , not of any particular numerical magnitude. A process can produce a fact and still fail to certify it metrically; a minimal closed process can certify if its closure structure is right.

Note the layered dependence: certification presupposes a fact to certify. The certification functional acts on processes that have already satisfied the $\beta_1 \geq 1$ condition and produced a trapped bit; it does not act on folds directly. This is what distinguishes level 3 of the hierarchy from level 2.

The functional

In the Planck Time companion paper, a fold process attempting to certify a realized non-redundant bit at an interface is constrained from two sides simultaneously:

Quantum distinguishability (QSL): $\Delta t \geq \kappa \hbar / E_{\text{dyn}}$

Gravitational admissibility (NTC): $E_{\text{grav}} \leq \eta(c^5/G)\Delta t$

connected by the conservative bridge inequality $E_{\text{dyn}} \leq E_{\text{grav}}$. The minimax over these two constraints yields:

$$\Delta t^2 \geq (\kappa/\eta)(\hbar G/c^5) \Rightarrow \Delta t \geq \sqrt{(\kappa/\eta)} \cdot t_{\text{P}}$$

up to order-one factors. The corresponding spatial form is $R_{\text{min}} \sim \ell_{\text{P}}$, which follows from Δt_{min} via the causal-interval identification $R \sim c\Delta t$ within the emergent metric (Planck Time paper, §9.3). The spatial and temporal forms are two representations of the same structural fixed point rather than independent results.

This is the certification functional. In structural language:

$$C[F] = (E_{\text{dyn}}(F) \cdot \Delta t(F)) / \hbar, \text{ subject to } E_{\text{grav}}(F) \leq \eta(c^5/G) \cdot \Delta t(F)$$

with the certification condition $C[F] \geq \kappa$. The Planck regime is the saturating case in which the bridge inequality becomes tight. The structural-language version and the physical-energy version are the same object viewed through different vocabularies; the structural version emphasizes that certification is a property of closure structure, the physical version makes the constraints concrete and the bound derivable.

What the integration gives

Three things directly. First, the Planck scale is derived, not stipulated: it is the smallest metric content certified by the minimal closed fold process that saturates the certification functional. Second, the prefactor $\sqrt{(\kappa/\eta)}$ is physically meaningful, depending on the choice of QSL form and the safety margin in the no-trapping condition. Third, the self-consistency check — below t_P the bit-certification rate cannot sustain the effective causal metric in which the constraints were stated — provides the boundary-of-validity argument that earlier drafts of this paper framed as a separate desideratum.

What remains open is narrower than earlier drafts suggested. It is no longer "construct the certification functional"; it is the substrate-level question identified in §12 of the Planck Time paper: derive QSL and NTC from substrate dynamics directly, without importing them from the effective causal description.

Interface with quantum mechanics

At the certification level, the QSL constraint reflects the minimal structure required for distinguishable evolution, providing the route by which the present hierarchy interfaces with quantum dynamics. The QSL is not a primitive of the substrate; it is the effective-level expression of the requirement that a certification process produce a distinguishable outcome, and it is what the certification functional imports from standard quantum mechanics. The present architecture does not replace quantum mechanics — it supplies the structural level at which quantum distinguishability first becomes the operative constraint, with the Planck scale marking the lower boundary below which distinguishability and admissibility cannot be jointly satisfied. Quantum mechanics and the certification functional are therefore not rivals; they are the same constraint viewed from different structural directions.

8. The coherence functional: integration with the Mesoscopic derivation

The closure scale ξ requires its own functional, distinct from the certification functional of §7, at a different structural level, answering a different question. The Mesoscopic Coherence companion paper supplies this object.

The functional

A fold-network structure of size L attempting to maintain intrinsic closure faces two competing failure modes:

UV failure (microscopic instability): $F_{UV}(L) \sim (\ell^*/L)^\alpha$ — fluctuations dominate at small L

IR failure (cosmological decoherence): $F_{IR}(L) \sim (L/L_{IR})^\beta$ — coherence loss dominates at large L

For $\alpha = \beta = 1$, which is forced (not merely generic) by smoothness and locality unless additional structure forbids it, the balance condition $F_{UV} = F_{IR}$ yields:

$$\xi \sim \sqrt{(\ell \cdot L_{IR})^*} \approx 30\text{--}100 \mu\text{m}$$

The Mesoscopic Coherence paper establishes that this result is overdetermined: four physically distinct mechanisms (instability minimization, leakage–mismatch balance, bandwidth matching, entropy-flux equilibrium) reduce to the same mathematical skeleton under minimal assumptions, and a fifth (information capacity crossover) arrives at the same scale through structurally independent mathematics.

Two functionals at two levels

The certification functional (§7) and the coherence functional (§8) share a family resemblance — both are minimax balances between competing failure modes at the UV and IR ends — but they are not the same object.

Functional	Level	Question	Variables	Failure modes balanced	Result
Certification	3	When does a realized fact acquire metric content?	\hbar, G, c	QSL vs. NTC (local)	ℓ_P
Coherence	4	When do facts form self-sustaining networks?	ℓ^*, L_{IR}	UV vs. IR (global)	ξ

They take different inputs, answer different questions, operate on different structural objects (single bit vs. fact network), and sit at different structural levels.

Interface with general relativity

At the coherence level, the density-to-geometry coupling (G) becomes manifest as effective curvature response, providing the route by which the present hierarchy interfaces with emergent spacetime descriptions. The coherence functional characterizes when fact networks first become self-sustaining; once they are, their committed density couples to effective geometric response via G , and the resulting structure is what standard treatments call the classical gravitational regime. The present architecture does not replace that regime — it supplies the structural level at which it first becomes available, with ξ marking the coherence threshold above which self-sustaining fact networks can carry the density-to-geometry coupling in a stable form.

9. Planck \neq ξ : a non-equivalence proposition

A referee will ask: could the Planck scale and the coherence scale be two ways of describing the same threshold, distinguished only by how it is analysed? This section shows directly that they cannot.

Proposition (Non-equivalence)

Any functional that fixes a threshold satisfying both

(i) minimal certification — local, single-bit, constrained by the joint action of the quantum speed limit and gravitational admissibility, and

(ii) intrinsic closure — global, multi-bit network, constrained by the balance of microscopic instability against cosmological decoherence

must reduce to two distinct sub-functionals at two distinct structural levels. Equivalently: the Planck scale ℓ_P and the coherence scale ξ are formally distinct thresholds; no single threshold fixed by a single functional can play both roles.

Argument

The argument proceeds from the structural features of the two functionals, which are incompatible in three independently sufficient respects. A fourth respect (scaling) supports the conclusion more loosely and is discussed separately.

Different variables. The certification functional is defined over \hbar , G , and c — the three bridge constants characterizing the transition from realized fact to metric content. The coherence functional is defined over ℓ^* and L_{IR} — the UV and IR scales characterizing the transition from isolated realized facts to self-sustaining networks. A single functional cannot be simultaneously a function of (\hbar, G, c) alone and of (ℓ^*, L_{IR}) alone; these are different input sets at the level of the functional's definition, independently of how the inputs may later be related.

Different structural objects. The certification functional acts on a single certification region containing one realized bit; the minimax balances energy scales *within* that region. The coherence functional acts on a fact network spanning many bits; the balance is between microscopic boundary-leakage of the network and cosmological decoherence *across* the network. The functionals take different kinds of object as input — a single process vs. a spatial extent of many processes — and no single threshold can characterize both.

Different mathematical structure. The certification functional is a *local minimax* between dynamical energy (driving distinguishability up) and gravitating energy (constrained by admissibility), at a fixed certification region. The coherence functional is a *global geometric-mean balance* between UV-increasing and IR-increasing failure rates, over variable spatial

extent. A local minimax at a point and a global balance over extent are different mathematical operations; they cannot in general share a fixed point.

Any one of these three considerations suffices to establish non-equivalence. Taken together, they establish it structurally: the Planck scale and the coherence scale are not two descriptions of a single object but two thresholds fixed by distinct functionals over distinct variables.

A fourth, weaker consideration (scaling). It is tempting to add that ℓ_P scales with (\hbar, G, c) while ξ scales with (ℓ^*, L_{IR}) and that the two scalings are independent. This is true at the level of the functionals' own definitions, but the support it lends to non-equivalence is weaker than the three structural considerations above, because under full closure the UV input ℓ^* of the coherence functional may trace to ℓ_P . In that case ξ acquires a downstream dependence on (\hbar, G, c) through its ℓ^* input, and the scalings are not fully independent. The correct reading: the certification functional is defined without reference to L_{IR} , and the coherence functional is defined without reference to (\hbar, G, c) directly; the two functionals therefore take distinct input sets at the definitional level even if their inputs are ultimately related under closure. This is parameter-independence at the level of the functionals, not full scaling independence, and the non-equivalence conclusion rests on the first three considerations rather than this fourth one.

Consequence

A competing programme proposing a single threshold that collapses Planck and ξ must either (i) identify additional physics that makes the three structural differences above illusory, or (ii) pay the cost of conflating two genuinely different questions — metric admissibility of a single realized fact versus intrinsic coherence of fact networks — into a single answer.

The architecture laid out in this paper takes neither option: it keeps the two thresholds distinct, integrates their respective functionals at their proper structural levels, and lets the orders-of-magnitude separation stand as a feature rather than a puzzle.

10. Claimed, assumed, open

Claimed (structural — level 1). The fold is pre-metric. It has no intrinsic length, duration, or mass. Metric content is not built into the primitive.

Claimed (structural — level 2). The minimal fact — a topologically trapped, irreversible, realized bit — is a distinct structural level between fold and metric. It is enabled by the condition $\beta_1 \geq 1$ on the fold-interaction graph but is not identical with that condition: the condition is what enables fact formation; the fact is what the condition makes possible.

Claimed (metric — level 3). The Planck scale is the floor of metric admissibility for a realized fact, not the size of a fold. It arises at the bridge between realized facts and physical metric, through the constants c, \hbar, G . The dimensional uniqueness of the Planck composites follows from there being exactly three bridge constants.

Claimed (metric — level 4). The closure scale ξ is a higher threshold — the mesoscopic scale at which realized facts form self-sustaining networks. It is not the metric floor and not the fundamental unit. Fold, minimal fact, Planck scale, and ξ sit at four distinct structural levels and must be kept formally distinct.

Claimed (non-equivalence). The Planck scale and the closure scale are formally distinct thresholds. They are fixed by different functionals over different variables, act on different structural objects, and have different mathematical structure. No single threshold can play both roles.

Claimed (operational). Sub- ξ facts exist, but only through environment-mediated closure: they ride on the coherent ξ -scale and macroscopic structure that envelops them. Isolated facts can occur at any scale above the Planck floor; self-sustaining fact networks require ξ .

Claimed (integration). The certification functional that fixes the Planck threshold is supplied by the QSL+NTC minimax of the Planck Time companion paper. The coherence functional that fixes ξ is supplied by the UV/IR geometric-mean balance of the Mesoscopic Coherence companion paper. The structural content of this paper is not the construction of those functionals but the architectural frame that locates them at their proper levels.

Assumed (bridge). That c , \hbar , and G correctly capture the temporal, quantization, and relational aspects of fold dynamics at the bridge, at the effective level. This is an import, not a derivation. See §5 for explicit discussion of status and direction.

Assumed (regime). That the QSL and NTC apply within the effective causal regime. The Planck Time derivation argues this is a self-consistency condition — the constraints identify the boundary of their own validity — but a fully substrate-level treatment that bypasses the effective regime is not yet in hand.

Open (primary). Substrate-level derivations of QSL and NTC. Both constraints are imported into the certification functional from the effective causal description. A fully closed bootstrapping argument would derive these constraints (or their substrate analogues) directly from tick dynamics on the void, without reference to the emergent level. This is the principal open object of the programme as currently constituted.

Open (secondary). Whether the topological condition $\beta_1 \geq 1$ can be derived as a structural necessity from the fold primitive itself, rather than imposed as a separate condition on fold-interaction graphs. The distinction is between *conditional* and *unconditional* necessity: the Topological Threshold derivation establishes that $\beta_1 \geq 1$ is necessary for local fact formation *conditional on the fold-interaction graph being given*, which is delivered. What remains is to derive the graph structure itself from the fold primitive plus minimal closure assumptions — i.e., to show that $\beta_1 \geq 1$ is forced *unconditionally* from the primitive, not merely required once the graph is specified.

Open (deepest — bridge-constant derivation). Whether c , \hbar , and G can themselves be derived from fold dynamics rather than imported at the effective level. The partial result $c = (L_{IR}^2 \cdot \hbar G$

$/\xi^4)^{1/3}$ from the companion to the Mesoscopic Coherence paper provides one direction — c from \hbar , G , and scale inputs — but leaves \hbar and G as imports. It should be acknowledged that under the present architecture this relation is self-referential by design: ξ is derived from ℓ^* and L_{IR} , and ℓ^* may trace to $\ell_P = \sqrt{(\hbar G/c^3)}$ under full closure, making \hbar , G , $c \rightarrow \ell_P \rightarrow \ell^* \rightarrow \xi \rightarrow c$ a consistency relation rather than a derivation from independent inputs. The loop is already closed at the level of internally consistent effective relations, though not yet as a substrate-level derivation from the primitive. The open object is therefore not "close the loop" — the loop is already consistent, by design — but "derive c , \hbar , and G from substrate-level fold dynamics that do not presuppose any of them," which would upgrade the consistency closure into a derivation from independent primitives.

11. Conclusion

The fold explains why physics must be built from distinguishable committed structure. The minimal fact layer explains what the first realized structure is — a topologically trapped irreversible bit, produced when fold interactions satisfy $\beta_1 \geq 1$. The Planck scale explains why realized facts have a smallest meaningful metric floor, fixed by the certification functional. The closure scale ξ explains why self-sustaining fact networks appear only at a mesoscopic threshold, fixed by the coherence functional; sub- ξ facts exist only through environment-mediated coupling to that threshold and above.

These four are not alternative answers to the same question. They are answers to four different questions, each at its own structural level. The fold is the ontological minimum; the minimal fact is the first realized structure; the Planck scale is the metric minimum; the closure scale ξ is the coherence minimum. The non-equivalence of Planck and ξ is a formal structural result, not a convention: they are fixed by different functionals over different variables and cannot be collapsed into a single threshold.

The two metric thresholds are now fixed by explicit functionals constructed in companion papers — the certification functional of the Planck Time derivation and the coherence functional of the Mesoscopic derivation. The two structural levels (fold and minimal fact) are established by the programme's foundational work and the Topological Threshold derivation respectively. What this paper contributes is not the construction of those objects but the architectural frame that locates them at their proper levels, establishes their non-equivalence formally, and prevents the conflations that would otherwise damage the programme.

The remaining open object is narrower than it first appeared. It is no longer the construction of any of the functionals, which exist. It is the substrate-level derivation of the constraints — QSL, NTC, and ultimately the bridge constants c , \hbar , G themselves — bypassing the effective causal regime in which they were originally formulated. That is the next problem, and the architecture laid out here is what it will close against.

This layered structure also explains why physics exhibits quantum discreteness, relativistic propagation limits, and macroscopic coherence as distinct but connected phenomena: each

emerges at a different level of the same underlying hierarchy. Quantum discreteness is the certification-level expression of the \hbar aspect of fold commitment; relativistic propagation is the certification-level expression of the c aspect; macroscopic coherence is the work of the closure level, above which self-sustaining fact networks can carry density-to-geometry coupling via G . Three apparently separate features of physics turn out to be three manifestations of the same underlying architecture, each at the structural level where its associated bridge constant becomes operative.

References

Companion papers within the VERSF programme

Taylor, K. *The Minimal Asymmetry*. AIDA Institute / VERSF Theoretical Physics Programme. versf-eos.com.

Taylor, K. *Why the Fundamental Unit Must Be the Minimum Distinction*. AIDA Institute / VERSF Theoretical Physics Programme. versf-eos.com.

Taylor, K. *The Topological Threshold for Fact Formation*. AIDA Institute / VERSF Theoretical Physics Programme. versf-eos.com. — Establishes the $\beta_1 \geq 1$ condition for local fact formation and the structural derivation of the $k_B \ln 2$ entropy scale. Cited in §§4, 6, 10.

Taylor, K. *The Planck Time as a Minimum Bit-Certification Scale: A Derivation from Bit-Contrast Balance and Ticks-Per-Bit Principles*. AIDA Institute / VERSF Theoretical Physics Programme. versf-eos.com. — Constructs the certification functional via the QSL+NTC minimax; derives $\Delta t_{\min} \sim t_P$ and $R_{\min} \sim \ell_P$. Cited in §§5, 7, 10.

Taylor, K. *Multiple Structural Derivations of the Mesoscopic Coherence Scale ξ_{meso} : Overdetermination of the Two-Planck Window from Independent Closure Principles*. AIDA Institute / VERSF Theoretical Physics Programme. versf-eos.com. — Constructs the coherence functional via the UV/IR geometric-mean balance; five independent derivations yielding $\xi \sim \sqrt{\ell^* \cdot L_{\text{IR}}} \approx 30\text{--}100 \mu\text{m}$. Cited in §§6, 8, 10.

Taylor, K. *Testing the Mathematics: The Speed of Light as a Computational Throughput Limit*. AIDA Institute / VERSF Theoretical Physics Programme. versf-eos.com. — Derives the closure relation $c = (L_{\text{IR}}^2 \cdot \hbar G / \xi^4)^{1/3}$. Cited in §§5, 10.

Foundational external references

Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D*, 7(8), 2333–2346.

Bousso, R. (2002). The holographic principle. *Reviews of Modern Physics*, 74, 825–874.

- Hatcher, A. (2002). *Algebraic Topology*. Cambridge University Press. — Background for Betti numbers and homology used in the topological threshold argument.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191. — Origin of the $k_B \ln 2$ bound; reinterpreted structurally via topological trapping in the companion Topological Threshold derivation.
- Lloyd, S. (2000). Ultimate physical limits to computation. *Nature*, 406, 1047–1051. — Prior two-sided squeeze yielding Planck-scale bounds on computation; algebraic structure shared with the Planck Time companion paper, with differences in operational target and ontological framing discussed there.
- Mandelstam, L., & Tamm, I. (1945). The uncertainty relation between energy and time in non-relativistic quantum mechanics. *Journal of Physics (USSR)*, 9, 249–254. — Quantum speed limit used in the certification functional (§7).
- Margolus, N., & Levitin, L. B. (1998). The maximum speed of dynamical evolution. *Physica D*, 120, 188–195. — Alternative QSL formulation; equivalent \hbar/E scaling.
- Mead, C. A. (1964). Possible connection between gravitation and fundamental length. *Physical Review*, 135, B849–B862. — Early identification of a minimum measurable length from combined quantum-gravitational considerations.
- Ng, Y. J., & van Dam, H. (2003). Spacetime foam, holographic principle, and black hole quantum computers. *International Journal of Modern Physics A*, 20, 1328–1335.
- Penrose, R. (1996). On gravity's role in quantum state reduction. *General Relativity and Gravitation*, 28, 581–600. — Background for gravitational decoherence arguments related to the NTC.
- Thorne, K. S. (1972). Nonspherical gravitational collapse: A short review. In *Magic Without Magic*, ed. J. R. Klauder. Freeman. — Hoop conjecture; the NTC used in the Planck Time companion paper is deliberately weaker than the full hoop conjecture.
- Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In *Complexity, Entropy and the Physics of Information*, ed. W. H. Zurek, pp. 3–28. Addison-Wesley. — "It from bit" programme; the VERSF primitive identifies the minimum distinction with Wheeler's bit at the ontological level.
- Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75, 715–775. — Background for the decoherence-as-trap-formation interpretation used in the Topological Threshold paper and inherited here at levels 2 and 4.