A Guide to VERSF

The Void Energy–Regulated Space Framework: Principles, Layers, and How to Read the Programme

Keith Taylor — VERSF Theoretical Physics Programme

Why This Guide Exists

The VERSF research programme now spans over 140 papers. Each paper is written to stand on its own shoulders — it assumes what came before and builds what comes next. That works well for a reader who started at the beginning and followed along. It works badly for someone arriving for the first time.

This guide is the front door. It does not prove anything new. It explains what VERSF is trying to do, what it assumes, what it derives, and in what order the pieces fit together. A serious reader should be able to finish this guide in an afternoon and then know which specific papers to read next for whichever aspect of the framework they care about.

The guide opens with a table of contents, then two parallel walkthroughs telling the same story at different registers: *How Reality Emerges (For the General Reader)* in plain language, and *How Reality Emerges: A Structural Walkthrough* in technical language. Either gives the shape of the claim. A reader can start at the register they find comfortable and move to the other if they want. The guide proper then has six parts:

1. **What VERSF actually assumes** — the handful of starting commitments the whole programme rests on.
2. **The layer stack** — how the framework builds upward from the most primitive concept to observable physics.
3. **The spine of the argument** — the chain of reasoning that leads from "physics requires facts" to the specific mathematical structure called the fold.
4. **A glossary of the key terms** — because several words get used in closely related but distinct ways, and it saves everyone time to pin them down.
5. **Reading paths** — suggested sequences through the papers depending on what you want to understand.
6. **What is still open** — honest accounting in three tiers (established, partially derived, open), followed by specific technical exposures.

Table of Contents

How Reality Emerges (For the General Reader)

The rest of this guide describes what VERSF claims about reality in technical language. What follows tells the same story without the specialist vocabulary — not simplified, just translated. If you come away understanding this, you understand what VERSF is claiming. If you then want the technical version, read the structural walkthrough that follows.

One clarification first. VERSF doesn't say reality appeared "over time," as if there was a time before and a time after. The story below walks up a ladder of dependencies: each rung rests on the one below. But the whole ladder is there at once, not built stage by stage. Where this description says "first this, then that," it means "this supports that" — not "this happened before that."

There has to be something, and it has to be something with differences. For there to be a world at all, two things have to be the case. There has to be a substrate — something that sets limits, absorbs consequences, and provides the background against which anything can happen. And there has to be difference: something has to be different from something else, or there's nothing to talk about. VERSF calls the substrate "the Void." The word is misleading — the Void isn't empty; it's the active set of constraints that reality operates under. The other primitive, difference itself, VERSF calls distinguishability. Neither is prior to the other. Each needs the other to make sense.

Difference alone isn't enough. It has to be the right kind of difference. A world where things differ but the differences don't stick — where yesterday's difference is gone today — isn't a world physics can describe. Physics needs differences that are stable (they stay put once they happen), combinable (what distinguishes A from B can be put together with what distinguishes B from C to tell you something about A and C), and accessible (a suitable observer could in principle tell they exist). Call this the threshold. Below it, you have raw being without structure. Above it, you have the kind of being physics can be about.

Above the threshold, a specific structure has to exist. When you work out what it takes to have stable, combinable, accessible differences, a particular minimum structure is forced. VERSF calls it the fold. You can think of the fold as a one-way door. On one side are possibilities — things that could be different, but aren't yet fixed. On the other side are facts — things that are definitely the case. The fold is where possibility becomes fact. Everything else in physics is built on this.

The first fact is when something definite happens. Before the first commitment happens, nothing has actually occurred — there are just possibilities. A commitment event is the moment when an option stops being a possibility and becomes an actual fact, irreversibly. The first such event is the first real thing — the first piece of "this, not that" anywhere in reality. Every later fact depends on earlier ones; reality builds itself up out of accumulated commitments.

Time is what happens when commitments pile up in order. Time isn't built in at the start. Before any commitments, there is no physical time — there's a kind of raw sequencing in the pre-commitment dynamics, but it carries no arrow and nothing a clock could measure. Once commitments start happening in order, the ordering *is* time. And the reason time has an arrow — the reason it seems to flow forward and not backward — is that each commitment is one-way: when an alternative is discarded, it goes into the Void, and the Void doesn't give it back. Commitments therefore stack up in a direction that can't be undone, and that direction is what you experience as the flow of time. This differs from some textbook explanations of time's arrow that appeal to how the universe started; here the arrow is built into the structure of commitment itself, not inherited from initial conditions.

Facts accumulate, and their density does the work. As commitments happen, they build up a kind of density across the substrate. Individual commitments are too small for us to see directly, but the pattern they make at large scales is visible — and that pattern is what we call mass-energy density. What looks to us like "a heavy object sitting somewhere" is, underneath, a region with a high density of committed facts. Mass is concentrated commitment.

Space and gravity are the substrate's pushback. Here is the surprising part. Distance, duration, and curvature — the geometric content of the world — aren't built in from the start. They are what happens when the Void pushes back against density patterns. The substrate has finite capacity; crowded regions resist further commitment more than empty regions do. The way these resistance-gradients pattern themselves is what a finite observer experiences as distance, time, and gravitational attraction. Gravity isn't a force acting in pre-existing space; it's the substrate's response to uneven loading, and space is that response.

There are minimum scales below which the normal descriptions stop applying. Nothing smaller than about 10^{-35} metres — the Planck scale — can meaningfully be called a "distance." At scales that small, the question of size stops having an answer. And below about 10^{-4} metres — the coherence scale, roughly the width of a human hair — single facts can still occur, but they can't sustain themselves alone; they need surrounding structure to hold them in place. This is why atomic and subatomic physics works the way it does: small-scale facts are real, but they ride on the larger structures around them.

The substrate has a specific shape, pinned to a specific number. When you ask what the minimum structure is that can hold a stable piece of reality together, a specific integer appears: seven. It takes exactly seven independent conditions to commit a unit of reality. Fewer, and the structure is ambiguous — it could be too many different things at once. More, and no configuration fits them all. Seven is the balance point. In the papers this integer is called $K = 7$, and it turns out to be the same integer that controls several specific numbers in physics: how strong electromagnetism is, how small the cosmological constant is, roughly what a proton weighs, and others. These are not independent mysteries in VERSF — they are different measurements of the same underlying structural number.

From here, everything else follows. The four fundamental forces. The elementary particles of the Standard Model. The specific values of the physical constants. In VERSF these are not independently given — they are consequences of the structure above, expressed in different mathematical languages.

The world you see is what all of this produces. Atoms, gravity, time, space, causality, the laws of physics and the specific numbers that appear in those laws — none of them are built in. They are what happens when a reality with the structure above gets looked at by a finite observer.

That is the shape of the VERSF claim. The rest of the guide works it out in technical detail.

How Reality Emerges: A Structural Walkthrough

The general-reader walkthrough above gives this same picture without technical vocabulary. What follows is the technical rendering — the specialist version of the same story, in a single continuous reading, stating what VERSF claims reality is and how its successive levels depend on one another. This section is the *shape* of the claim — a walkthrough, not a derivation or an axiom list.

A note on tense and voice. What follows is structural, not historical. There is no moment at which the Void "existed alone" and then distinctions "appeared." VERSF's proto-time — the sequencing of reversible pre-commitment dynamics — is not physical time, and physical time itself is derived from commitment sequences several levels into the walkthrough. The emergence described here is a chain of *structural dependencies*: each level is what the level below makes possible. It is not a sequence of events unfolding against a pre-given temporal background. Where English makes that distinction hard to carry, the language leans on "underwrites," "rests on," and "makes possible" rather than on "then" or "next."

The walkthrough has ten stations, corresponding broadly to the guide's layer stack but organised by dependency rather than by formal layer-number. Readers who prefer navigation by physics domain should consult Part 2's compressed vertical view; readers who want the formal derivation should consult Part 3's twelve-step spine and five-step admissibility chain. This section gives neither — it gives the shape.

The sub-structural pole. At the base of the architecture are two co-primitives: the Void and distinguishability. The Void is the constraint medium — that which enforces finite capacity on bounded regions, absorbs discarded alternatives, and resists over-concentration of commitment. Distinguishability is the structure of actualisation — that which enables distinctions to be made, at finite resolution, stably. Neither is prior. The Void without distinguishability is indistinguishable from nothing; distinguishability without the Void has nothing to actualise from. They are complementary in the same structural sense as wave and particle in Bohr's formulation: each required for the intelligibility of the other, neither reducible to the other, no third description encompassing both. Everything else in the framework is what this complementarity makes possible.

Distinction as floor. Distinguishability-bearing structure is the floor of physics. A reality in which nothing differs from anything else is structurally indistinguishable from no reality at all. Below this floor there is, in the formal sense developed in *Distinction as an Ontological Precondition for Structural Reality*, no object that physics can be about. The floor is a *necessary* condition on any reality with structure; it is not yet *sufficient* for physics. A reality can be distinction-bearing without being able to support stable, reproducible, compositional facts.

Admissibility as threshold. Sufficiency arrives when distinction-bearing structure additionally satisfies the admissibility conditions: finite distinguishability (A1), structure-independent comparison (A2), observational accessibility (A3). At this threshold, distinction becomes capable of supporting physics — of producing facts that persist, compose, and can be referred to. The floor-to-threshold gap is the subject of the in-preparation successor paper *Ontological Bookkeeping of the Admissibility Axioms*.

The fold. Once a reality is admissibility-capable, it contains a specific minimal structure: the fold. The fold is a two-dimensional intrinsic commitment interface, topologically a sphere, separating the reversible pre-commitment sector from the irreversible committed sector. It carries two binary geometric degrees of freedom — commitment polarity $\sigma \in \{0,1\}$ and reversible orientation $\omega \in \{-1,+1\}$ — and no more. The pre-commitment dynamics on the fold are unitary evolution on $\mathbb{C}^4$: the dimension comes from the four-state structure (σ, ω), and the complex numbers are forced by admissibility acting on the classification of normed division algebras. The fold is not chosen. It is the minimal structure that admissibility permits. Nothing smaller supports facts; nothing larger is minimal.

The first realised structure. The fold makes commitment *possible* but is not itself a commitment. A commitment event occurs when the fold-interaction graph supports topological trapping — the condition $\beta_1 \geq 1$, under which alternative routes permit one alternative to enter a loop while the other does not, producing a non-recombinable distinction. The first trapped bit is the minimal fact: the first irreversible realised structure in the architecture. It is pre-metric — it has no associated length, duration, or mass at this level — but it is real. Something has definitely occurred. Given reversible microdynamics and local construction, the trapping condition is not merely permissive but *productive*: wherever $\beta_1 \geq 1$ is satisfied, minimal facts generically form rather than merely being possible. The minimal fact is also the first unit of entropy (k_B ln 2, per Landauer) and the first increment from which the arrow of time is later constructed.

Physical time and its arrow. The first commitment event is also the first step in constructing physical time. Proto-time — the reversible pre-physical ordering on the fold — carries no arrow and no metric; it is the sequencing the pre-commitment sector has when it has dynamics but not yet facts. Physical time arises when commitment events start accumulating in ordered sequences, and its arrow is the direction of commitment itself. Each commitment externalises its discarded alternatives to the Void, and the Void absorbs but does not return; the sequence of commitments is therefore ordered in a direction that cannot be reversed without undoing absorbed distinguishability, which the Void's capacity structure does not permit. The arrow of time is therefore architectural, not accidental — not inherited from low-entropy initial conditions, not dependent on a Past Hypothesis, not a statistical coincidence that might have gone the other way. It is the direction commitment runs. At scales large compared to ξ, the partial-order structure of commitment events recovers the causal ordering of special relativity; combined with the invariant propagation speed derived at Layer 5, this yields emergent Lorentz invariance at the macroscopic level.

The record and the entropy field. Accumulated minimal facts constitute the record field $\rho(x,t)$, which tracks the density of committed distinctions across the substrate. Coarse-graining ρ over causal neighbourhoods at the fundamental statistical scale R^* produces the smooth entropy field $s(x)$. ρ is the *load* the Void is carrying in each region; s is the coarse-grained signature of that load. The scale R^* is a statistical-resolution parameter, not a spatial direction — coarse-graining is re-description at finite resolution, not motion into a new dimension; see Layer 5 on *Depth Is Not a Spatial Direction* for why this distinction matters. The Single-Source Theorem establishes that all observable physics is ultimately a functional of ρ and its derivatives — no hidden degree of freedom survives outside this field, because any such degree of freedom would either be observable (and hence already captured in the fold's observable algebra) or unobservable (and hence eliminated by admissibility).

Geometry as Void response, and the metric floor. The Void responds to gradients in s(x) by modulating the rate and direction of commitment propagation. What finite observers register as distance, duration, and curvature is this response, expressed in the continuum limit as an effective pseudo-Riemannian metric $g_{\mu\nu}(x)$ built from s and its gradients. Geometry is not imposed on the framework and not derived from matter on a pre-existing manifold. It is the Void's structural response to entropy loading — a response field, not the smoothed field. Metric content first becomes *certifiable* at the Planck scale, which is the floor at which a realised fact can acquire operationally meaningful length, duration, and mass. This floor is fixed by the joint constraint of quantum distinguishability and gravitational admissibility. Below the Planck floor, facts can still form but metric questions about them cease to have answers.

Coherence and the mesoscopic threshold. A single minimal fact, or a handful, can occur above the Planck floor without forming a self-sustaining network. Networks of facts — structures that hold together without external scaffolding — require the coherence threshold $\xi \approx 8 \times 10^{-5}$ m, the mesoscopic scale at which the balance between UV instability and IR decoherence becomes favourable. Above ξ, commitment structure is *intrinsically* closed: it sustains itself. Below ξ, facts still form, but only through environment-mediated closure — they ride on the ξ-scale surroundings that supply the global connectivity their local neighbourhoods cannot provide. This is why sub-ξ physics exists. Atomic-scale facts are real, but their existence as persistent records depends on the ξ-scale environment in which they are embedded. The Planck floor and the coherence threshold are formally distinct: they are fixed by different functionals over different variables and cannot be collapsed into a single scale.

The substrate and the downstream physics. Closure at and above ξ is not arbitrary. It is supported by a specific substrate architecture: the K=7 simplicial structure, forced by the No-Go theorem as the minimum cell count consistent with non-degenerate closure, fact-supporting topology, automorphism saturation, and minimal architecture. From K=7, the rest of observable physics follows. The κ-field emerges, with its mass coefficient $\sqrt{(4/3)}$ fixed by PGL(3,2) irreducibility acting on the non-uniform subspace V_6. The Standard Model gauge content emerges from hexagonal closure geometry. The particle spectrum emerges from fold-boundary closure modes. The four fundamental forces correspond to four distinct classes of closure relation. Specific numerical constants — the fine-structure constant with bare term $\alpha^{-1}_{bare} \approx 137.143$ derived from K=7 geometric constraints (the second-order correction to the observed 137.036 is in progress, with its structural inputs established in *Completing the Interface Bridge*), the cosmological constant Λ placed uniquely in the observed band via the ξ(K) coherence-scale derivation (absolute normalisation pending), the tensor-to-scalar ratio $r \approx 0.027$–0.033 from the κ-field spectrum, the area-linear scaling of Bekenstein–Hawking entropy with K = 7 in the coefficient (the exact 1/4 factor pending normalisation) — are geometric consequences of the closure architecture, not parameters fitted to observation. Layer 8 and Part 6 state precisely which of these are fully derived, which are partially derived, and which are open.

The macroscopic world. What we recognise as physics — atoms, fields, gravity, time, space, causality, the specific laws observed in the lab — sits at the top of this dependency chain. Every feature of it rests on the structure walked through above. Space is the Void's response to spatial entropy gradients; time is the ordering of commitment events; causality is the directed structure of those orderings; mass is commitment density; curvature is density pushback; the quantum mechanics of pre-commitment dynamics is $\mathbb{C}^4$ unitary evolution on the fold; the Standard Model gauge structure is K=7 hexagonal closure; the specific constants are geometric consequences of the closure architecture.

The walkthrough ends where physics begins.

What this section is not. It is not a derivation — Part 3's twelve-step spine and five-step admissibility chain do that work. It is not a list of assumptions — Part 1 gives those. It is not a layer-by-layer expansion — Part 2 gives those with full detail and paper references. It is the shape of the claim, in structural-dependency order, for a reader who wants to hold the whole picture before checking any of it.

What the guide does with this shape from here. Part 1 states what VERSF assumes and clarifies the status of those assumptions. Part 2 expands each level of the walkthrough with technical content and paper pointers. Part 3 gives the formal argument that each level is forced rather than chosen. Part 4 defines the terminology. Part 5 offers reading paths through the papers. Part 6 accounts honestly for what remains open. The rest of the guide is what this walkthrough requires if the shape is right.

Part 1: What VERSF Actually Assumes

Most physical theories start by assuming a lot. They assume spacetime exists. They assume there are fields living on it. They assume certain symmetries, certain equations, certain fundamental constants. The theory is then the study of what those assumptions imply.

VERSF takes a different approach. It starts from almost nothing and tries to show that a surprising amount of structure is forced by the very idea of "doing physics." The starting commitments are these:

1. Physics requires facts. Not facts in the everyday sense — facts in a precise technical sense. A physical fact is an outcome that is stable (it persists), recoverable (a finite observer can access it later), reproducible (different observers can identify the same result), irreversible (once it happens, a local observer cannot undo it), and distinguishable (it is separated from alternative outcomes by at least some minimum resolution).

Without facts of this kind, there are no experiments, no measurements, no data, no testable predictions. A "physical theory" that produces no facts is not physics — it is mathematics with no connection to observation. This is not a philosophical claim; it is a definition. Call this the *fact requirement*.

2. Finite distinguishability. In any bounded region of the universe, only finitely many distinctions can be stably recorded. You cannot fit infinitely many distinguishable outcomes into a finite volume. There is a minimum resolution below which distinctions simply cannot be made stable. This is what keeps physics from requiring infinite precision, which no finite recording system could ever supply.

3. Local causal boundedness. No finite observer has access to everything. An observer can only interact with a bounded region; they cannot simultaneously control every degree of freedom that influences what happens near them. This is the statement that locality, in some effective form, is a feature of physical observers rather than an accident.

4. Admissibility of description. If two descriptions of reality produce the exact same facts under every admissible protocol, they describe the same physical situation. Unobservable distinctions carry no physical content.

This is sometimes dismissed as an epistemic rule — a convention about how we talk about theories — rather than a physical constraint. But it is physical, and the argument runs as follows. A distinction that never registers as or contributes to any fact, and never influences the distribution of any future fact, has no handle through which a finite physical system could detect, propagate, or stabilise it. Under the other commitments — irreversibility (which forces physical processes to terminate in facts or in reversible evolution that eventually does), and finite distinguishability (which bounds how much uncommitted structure a bounded region can sustain) — any putative causal influence must eventually either commit to a fact or dissipate below the resolution threshold. A distinction that does neither is causally inert: it cannot affect, and cannot be affected by, anything that physically happens. A causally inert structure is not physics; it is surplus description. Admissibility is therefore not a rule about what we can know; it is a rule about what can have physical consequences, given the other three commitments. It is not an independent axiom so much as a consistency condition that falls out when finite observers, finite distinguishability, and irreversibility are taken together.

A note on observers. The word "recoverable" in the fact requirement above sometimes raises an objection: aren't you smuggling observers into the foundations of physics? The answer is no. The fact requirement is about recoverability in principle — whether any finite physical system within the causal reach could access the outcome — not about whether any actual observer is watching. A universe in which a fact occurs is a universe in which that fact *could* be accessed by a suitably situated finite system, regardless of whether such a system is there. This is a structural condition on physical processes, not a condition on consciousness, measurement apparatus, or human knowledge.

A point worth stating explicitly: in VERSF, commitments are *non-relational facts*. A commitment event is a fact *simpliciter*, not a fact relative to any particular system or observer. This distinguishes VERSF from Rovelli's relational quantum mechanics, in which every fact is fact-relative-to-a-system and no commitment can stand independently of the system to which it is relativised. In VERSF the distinction is not a matter of framing: the programme affirms a non-relational substrate (the Void) and holds that the facts it produces are themselves non-relational. Finite observers in the sense used above are the bounded subsystems *within which* facts can be recovered — they are not what *constitutes* facts. The distinction matters for foundational positioning; it does not affect the structural derivations.

Throughout the guide, the phrase "finite observer" is shorthand for "any bounded finite physical subsystem within the causal reach capable of stably recording distinctions, whether or not actually instantiated." It is an abstract structural notion, not a reference to a sentient agent or a measurement device. Similarly, "recoverable" and "reproducible" are properties of outcomes relative to such hypothetical bounded subsystems — the outcome *can be* recovered or reproduced, not the outcome *must be observed* by someone. The guide uses the shorter term for readability, but the longer meaning is always what's intended. Wherever "observer-dependent" concerns arise in what follows, this shorthand should be mentally expanded.

That is it. Everything else in VERSF — the fold, the commitment interface, $\mathbb{C}^4$, K=7, emergent time, emergent space, the specific forms of the laws — is claimed to follow from these four commitments plus careful reasoning.

A note on what these commitments are. They are not arbitrary axioms, and they are not pulled from philosophy. They are the operational face of something deeper: the constraint structure of the Void (introduced in Layer 0 of Part 2). The Void is the medium that enforces irreversibility, absorbs discarded alternatives, and bounds the capacity of bounded regions to hold distinguishable structure. The four commitments are what a finite observer sees when the Void's constraints act on the pre-commitment sector. You can read the programme two ways: top-down, starting from the four commitments and deriving the fold and everything downstream (which is how the main Uniqueness paper proceeds); or bottom-up, starting from the Void as constraint medium and showing that the four commitments are what its structure necessarily imposes on any finite observer. Both readings end in the same place. Part 1 uses the top-down reading because it is formally cleaner; Part 2 introduces the bottom-up reading because it is conceptually more illuminating.

Whether either reading convinces you is what the 140 papers are attempting to establish. Whether you find the four commitments plausible even before you check the details is largely a matter of how you feel about them. If you think they are unreasonably strong, the programme will feel over-reaching. If you think they are almost trivially weak — just what "doing physics" means — then the programme's results will feel genuinely surprising.

Most of the framework's appeal lies in taking commitments that feel weak and showing they force structure that feels strong.

A note on the three-level foundation. The four commitments above are a *front-door* framing: they list, in a single register, conditions that any fact-producing physics must meet. The formal programme organises the same content into three distinct levels, and a reader going from this guide to the foundational papers will encounter them in that form. The ontic level concerns what is foundational in being — and the current programme (see *The Layered Foundation of VERSF: Distinguishability, Potentiality, and the Architecture of Observation*) takes the ontic foundation to have *two* complementary primitives: potentiality (the Void) and actualisation (distinguishability realised as irreversible commitment), neither reducible to the other. The epistemic level concerns what is foundational for observation, and at this level distinguishability is primitive: nothing enters the space of data without being distinguishable. The methodological level concerns what is foundational for physics: the structural conditions on distinguishability that any theory of the observable must respect. The four starting commitments of this guide span these three levels without separating them — commitment 1 (facts) and commitment 4 (admissibility) are methodological; commitment 2 (finite distinguishability) is epistemic; commitment 3 (local causal boundedness) is both epistemic and methodological; and the Void that underwrites them all is ontic. This guide keeps the four-commitment framing because it is easier to hold in one reading; the three-level structure is what the papers work in.

A further refinement the current programme makes explicit is the distinction between the *floor* and the *threshold* of physics. The floor is distinction-bearing structure: the minimum condition for there to be a world with any structure at all, established in *Distinction as an Ontological Precondition for Structural Reality*. The threshold is admissibility: the further conditions a distinction-bearing structure must satisfy to support stable, compositional, accumulating records — that is, facts in the programme's technical sense. A reality can be distinction-bearing without being physics-capable; a reality cannot be physics-capable without being distinction-bearing. The ontological chain is: structurally-meaningful-reality (the floor) → physics-capable-structure (the threshold) → closure → uniqueness → the fold. The guide compresses these tiers because a front-door reading does not need them separated, but the papers keep them strictly distinct.

A note on axiom labelling. Where this guide uses the designations A1, A2, A3, these refer to the canonical admissibility axioms of the current foundational papers: A1 = finite distinguishability (a finite procedure partitions the candidate configurations into at least two classes), A2 = structure-independent comparison (comparison relations whose truth values do not depend on representation), A3 = observational accessibility (at least two outcomes of A1 are accessible to some finite observational procedure). Readers who dip into earlier derivation-chain papers will occasionally see A1–A3 used to label a different (but related) axiom set drawn from the specific derivation at hand; where this occurs in Part 6's discussion of open technical items, the labelling follows the source paper, and the guide flags the variation in context. The canonical labels above are the ones the guide uses when speaking for the programme as a whole.

A note on entry points. The framework admits two equivalent starting points, and this is not a casual feature — it is a deliberate structural choice of the programme. The first entry point is a meta-principle reconstruction based on facthood and admissibility: you accept the four operational commitments above (sometimes expanded as meta-principles M1–M5 in the reconstruction papers), and you derive the fold and everything downstream from them through closure, uniqueness, and saturation theorems. The second entry point is an ontological interpretation in terms of Void and commitment dynamics: you accept that reality contains a constraint medium that enforces finite capacity and externalises discarded alternatives, and you derive the operational commitments as the observable face of that structure. Both routes produce the same architecture. Both are valid. The guide uses both because they illuminate different aspects — the first is formally cleaner, the second is physically more suggestive — but a reader should understand that accepting one route does not require accepting the other, and the formal results do not depend on which entry point is preferred.

Part 2: The Layer Stack

VERSF is a layered framework. Each layer is built from the one below it. If you understand the stack, you understand where any individual paper fits. Here it is, from the most primitive concept at the bottom to observable physics at the top.

Before the detailed stack, a compressed vertical view of the full architecture, for readers who want to see the whole shape at a glance. This view organises the programme by physics domain; for a complementary view organised by structural dependency, see the opening walkthrough (*How Reality Emerges*). The two are meant to be read together: the dependency view shows why each level is what the one below makes possible; the domain view shows what physics each level contributes to.

Ontology. Commitment events and the record field $\rho(x,t)$. The Void as constraint medium.

Field structure. The κ-field as a fluctuation of ρ; the derived Lagrangian governing commitment-capacity dynamics.

Statistical mechanics. Poisson statistics of commitment events; the spectral density $J(\omega)$ of the commitment-event bath.

Memory physics. The memory kernel $K(\tau) \sim 1/\tau$, with the scaling and form derived from the statistical structure of commitment events (though a geometric ansatz enters the full derivation — see Part 6).

Dynamics. Nonlocal equations of motion for κ and ρ, following from the memory kernel and closure constraints.

Cosmology. The Friedmann equation, slow-roll conditions, the scalar spectral index, and the tensor-to-scalar ratio $r \approx 0.027$–0.033.

Emergence. Gravity as the Void's response to commitment-density gradients; the Born rule $|\psi|^2$ as the operational measurement rule for fold commitment outcomes.

This vertical structure corresponds to the nine horizontal layers below (Void, proto-time, fold, $\mathbb{C}^4$, $\rho(x,t)$, emergent time/space/causality, K=7 substrate, fields/forces, and predictions), but organised by physics-domain rather than by layer-of-abstraction. A reader who prefers to navigate by physics domain (cosmology, quantum mechanics, gravity) may find the vertical view more useful; a reader who wants to follow the structural derivation in order should use the detailed layer stack below.

Layer 0: The Void — Constraint Medium, Not Empty Substrate

The word "Void" can sound like absence or metaphysical baggage. In VERSF it is the opposite. The Void is not nothing — it is the active constraint medium that enforces the limits under which physical reality can exist. A reader who leaves this section thinking the Void is a placeholder will miss what the framework is actually claiming.

Operationally, the Void is the pre-commitment sector: the domain in which distinctions may exist reversibly but have not yet become facts. But the central claim is stronger: the Void is the medium that receives discarded alternatives and enforces the constraints of irreversible commitment. Its existence is not posited; Theorem 6.2a of *Structural Uniqueness of Physical Law from Fact Formation Constraints* proves irreversibility requires externalisation of information to physically real degrees of freedom inaccessible to bounded observers.

This gives the Void three active roles:

Absorption. It carries away the information about alternatives that are not realised. Every commitment event transfers the distinguishability of discarded outcomes into the Void, where it remains physically instantiated but operationally inaccessible to any finite observer.

Capacity enforcement. It limits how much distinguishable structure can be stabilised in any region. Bounded regions support finitely many facts because the Void imposes a capacity ceiling on commitment density.

Response. It resists over-concentration of commitment. Where local commitment density approaches the Void's capacity, further commitment becomes costly, and that cost manifests as structural resistance. What we recognise as gravity is this resistance expressed as geometric response — mass is commitment density, curvature is the Void's pushback against gradients in that density, gravitational attraction is the system relaxing toward more uniform loading. Geometry is the signature of how the Void deforms under load, not the Void itself.

The three roles are not independent features but expressions of a single constraint: the Void enforces the finite capacity of reality to hold distinguishable structure. The fact requirement, finite distinguishability, local causal boundedness, and admissibility are the operational face of this constraint structure — they are what finite observers see when the Void's constraints act on the pre-commitment sector. This is why the four starting commitments of Part 1 are not arbitrary axioms but are forced.

A note on primacy. The guide emphasises the Void as the active constraint medium, but the formal framework does not treat the Void as ontologically prior to distinguishability. The Void and distinguishability are *complementary* primitives in a precise sense developed in *The Layered Foundation of VERSF*: each is required for the intelligibility of the other, neither is reducible to the other, and no describable unification into a common substrate is available within any theory of the observable. The structural model for the claim is Bohr's complementarity in quantum mechanics — two descriptions, mutually required, with no third view encompassing both — used here by analogy rather than by inheritance. The Void provides the substrate of possibility (where alternatives can exist, where discarded ones can go); distinguishability provides the structure of actualisation (which distinctions become facts, at what resolution). Potentiality without actualisation is indistinguishable from nothing; actualisation without a substrate has nothing to actualise from. The two are co-primitive.

The Void's substrate role is itself twofold. It is the medium across which *past commitments persist*, enabling records to be referenced and composed with later commitments — without a persisting substrate, a purely relational ontology cannot support the compositional closure that physics requires. It is also the *source of pre-actualised possibility* from which new commitments are drawn — without a field of possibility, a framework that accommodates historical records has no account of how new facts enter the record. These two roles are distinct: the first is forced by the relational-closure argument (facts must persist for physics to be a science of relations among facts); the second is forced by the independent intelligibility of actualisation (commitments must be commitments *from* something). VERSF unifies the two roles in a single substrate as an economy, not a stipulation: one Void doing the work of past-persistence and future-possibility is more parsimonious than two, and it is the ontology the programme adopts because nothing weaker supports both roles.

The foundational architecture papers develop this complementarity explicitly. The guide's Void-forward framing is a pedagogical choice, not an ontological reduction. A reader who prefers a distinguishability-first framing will arrive at the same architecture by the same chain of consequences; a reader who wants the ontic account in full should read *The Layered Foundation of VERSF* and *Distinction as an Ontological Precondition for Structural Reality* together.

What the Void Is and Is Not

Because the Void does several kinds of structural work at once, it is easy to read it as an all-purpose mechanism — a "Void of the gaps" invoked wherever the framework needs a constraint. This section pins down what the Void is as a formal object by saying what it is not, what it is, and in what sense it is a single thing despite doing multiple jobs.

The Void is not a field. It carries no field operators in the QFT sense, no Lagrangian density in the standard sense, no equations of motion defined on a background manifold. Calling it a "field" would import a mathematical structure (a section of a bundle over spacetime) that presupposes the spacetime VERSF is trying to derive. Fields live on backgrounds; the Void is what geometry emerges *from*, not something that emerges *on* a geometry.

The Void is not a background. Backgrounds are passive — they sit there and provide a stage on which dynamics happen. The Void is active: it responds to load by producing the geometric effects we recognise as curvature and gravity. A passive background does no work; the Void does work.

The Void is not a manifold. Manifolds have points, neighbourhoods, differential structure, and metric relations. The Void has none of these as primitives. What manifold-like structure emerges in the large-scale limit is derived in Layer 5 from coarse-grained commitment statistics, not posited.

The Void is a constraint medium. This is the primary formal characterisation. It enforces finite capacity on bounded regions (a finite region cannot stabilise infinitely many distinctions), enforces irreversibility (discarded alternatives must be externalised somewhere), and enforces capacity saturation (commitment becomes costly as local load increases toward capacity). Constraints, not degrees of freedom, are what the Void carries.

The Void is an information-theoretic reservoir. This is the complementary characterisation: it is the destination for discarded distinguishability and the source of pre-actualised possibility. A reservoir is characterised by its capacity, its coupling to the system it receives from or supplies, and its response to loading. The Void has all three, which is why the reservoir description captures it at the operational level.

The Void is the ontic complement to distinguishability. This is the foundational characterisation: not a structure alongside distinguishability but its co-primitive in the complementarity sense developed above.

The three active characterisations — constraint medium, information-theoretic reservoir, ontic complement — are not three separate Voids in tension. They are three descriptions of the same structure, each appropriate to a different question. Asked "why is local capacity finite?", the answer is the constraint-medium characterisation. Asked "where do discarded alternatives go?", the answer is the reservoir characterisation. Asked "why is there commitment at all?", the answer is the ontic-complement characterisation. The Void's apparent multi-functionality is the different aspects of a single structure responding to different questions, not a list of separate jobs bundled into one term.

One diagnostic separates the Void from look-alike constructions in other frameworks. A "hidden sector" in a QFT-style theory would be a collection of additional fields on the same spacetime, observationally inaccessible but formally homogeneous with the observable sector. The Void is not that: it is not on spacetime, not composed of fields, and not formally homogeneous with the commitment sector. A "reservoir" in thermodynamics is typically characterised by its temperature, a finite-dimensional parameter. The Void's capacity is not a single parameter but a structural constraint on how much distinguishability a bounded region can hold, which is local rather than global. Calling the Void either "a hidden sector" or "a reservoir" in the narrow senses of those terms is a category error.

One further concern deserves naming directly: the *over-centralisation* diagnostic. A sophisticated critic could grant all of the above — that the Void is not vague, not a catch-all, formally distinct from hidden sectors and thermodynamic reservoirs — and still press a sharper objection. Much of the programme's derivational weight is currently carried by invocations of "the Void's response function": geometry is the Void's response to entropy gradients, gravitational strength is a parameter of the Void's response, the effective metric coefficients A(s) and B(s) are functionals of the Void's response, the constraints enforcing finite capacity are the Void's response to loading. Each of these invocations is structurally correct — the Void is indeed what all of these things respond to — but the response function itself is, at the current state of the programme, not fully derived. The risk the programme faces here is not that the Void is vague, but that the Void becomes the place where unsolved derivation tasks accumulate.

Acknowledging this risk and naming the specific papers that need to close it is the correct response; denying the risk would be overclaim. The principal items the Void's response function has not yet closed are named in Part 6: the explicit derivation of A(s) and B(s), the full linearised-GR reduction, the substrate-level derivation of the gravitational efficiency factor λ, and the full κ-field integration with the main entropy/geometry field. Each of these is a place where "the Void responds" currently stands in for "the Void responds *in the following explicit way*." Closing those derivations is how the programme converts structurally-correct invocations into derivations.

It is worth noting that the over-centralisation concern is not static — it is actively shrinking as specific Void-response behaviours acquire dynamical realisations. The gravity papers establish that the Void's response to commitment-density gradients manifests concretely as constraint-driven flux, continuity equations, and transport-limited propagation — not as abstract "response" invocations. The Poisson equation in the weak-gradient limit is derived, not asserted; the inverse-square law emerges as the unique admissible vacuum solution; the coupling G reduces to a single structural normalisation $(\lambda/\mathcal{C})\cdot\xi^2c^3/\hbar$ in which all geometric structure is fixed and only the efficiency factor λ remains open. Each of these results moves content *out* of the "Void responds mysteriously" category and *into* the "Void responds as a specific derived dynamical law" category. The Void is, increasingly, operationally realised rather than philosophically invoked. What remains philosophical in the Void's role is narrower now than it was, and the remaining items are specifically the ones listed in Part 6.

The relationship between the record field $\rho(x,t)$ (introduced in Layer 4) and the Void is worth flagging now: $\rho(x,t)$ is not just the distribution of committed facts but the *load* the Void is currently carrying in each region. When Layer 4 says all of physics is a functional of ρ, the deeper claim is that all of physics is a functional of how the Void is loaded.

Papers: *The Void Is Not Nothing*, *The Void Before Everything: Why the Simplest Pre-Cosmic State Is Pure Potential*, *Why Reality Needs a Void to Produce Facts*, *Distinguishability and the Void*, *Gravity as Critical Entropic Back-Pressure* (Void response producing gravity), *Void Anchoring: A Different Way to See Structure*.

Why the Void Cannot Be Eliminated

The Void's functional role — absorbing discarded alternatives, enforcing finite capacity, responding to load — is unavoidable for any framework that supports finite observers and irreversible facts. Any candidate replacement would have to allow irreversible commitment without externalisation, or permit unbounded local capacity, or both. Neither option is consistent with Part 1's commitments: without externalisation, discarded alternatives remain accessible and commitment is reversible after all; with unbounded capacity, bounded regions would need to stabilise infinitely many distinctions, which no finite physical subsystem can do.

One subtlety worth naming. The argument above rules out frameworks that lack a Void function but does not force the Void to be a single unified medium. A distributed substrate — many local Voids, or a structure in which each region has its own externalisation destination — would also satisfy the operational requirements. The claim is therefore narrower than "VERSF's Void is the unique substrate"; it is "the Void's functional role is unavoidable." Whether the implementation is singular or distributed is a separate question, addressed in the specific papers on Void structure and geometry, and relates to the ambient substrate dimension question listed in Part 6 as open. What cannot be eliminated is the function; the geometry of the implementation is a further question.

With that qualification: remove the Void function and you remove externalisation; remove externalisation and you remove irreversibility; remove irreversibility and you remove facts; remove facts and you remove physics. The Void is not an element of the framework bolted on at the bottom. It is the reason the framework has any structure at all. Any theory that produces physics contains a Void function, whether or not it names one.

Layer 1: Proto-Time

Proto-time is the pre-physical sequencing of the reversible dynamics that happen before a commitment event. It is not time as we experience it — you cannot put a clock against it, you cannot measure it, you cannot be in it. Proto-time is the logical ordering that governs how the reversible, pre-factual dynamics evolve before a fact crystallises.

Why does this matter? Because physical time, in VERSF, is *emergent*. It is generated by the ordered sequence of irreversible commitment events. Before any fact has occurred, there is no physical time. But there is still some notion of "what happens first" in the reversible sector, and that is what proto-time captures.

Papers: *Two Kinds of Time: Proto-Time and Physical Time*, *Where Time Really Comes From*, *Time Did Not Begin*.

Layer 2: The Commitment Interface (The Fold)

This is the central object of VERSF. The fold is the boundary between the reversible pre-commitment sector and the irreversible committed sector. It is where the crucial transition happens: where an unresolved alternative becomes a fact.

The fold is proved to be a two-dimensional surface — specifically, topologically equivalent to the surface of a sphere (S^2), as established in *Deriving the Minimal Commitment Structure at the Void–Universe Boundary*. It carries exactly two binary geometric properties: a commitment polarity (which side of the interface is committed, which is reversible) and a reversible orientation (the direction of traversal). These four states form the minimal structure needed to encode a commitment event.

The fold is not something imposed on the framework. It is the unique minimal structure, up to observational equivalence and field redefinition, that the four starting commitments force into existence. This is the core claim of the Uniqueness Programme.

Papers: *Deriving the Minimal Commitment Structure at the Void–Universe Boundary*, *The Fold: A Boundary Theory of Space, Quantum Correlations, and Gauge Symmetry*, *The Fold Interface Law*, *The Geometry of Facts: Why Reality Needs Two Dimensions First*.

Layer 3: The Pre-Commitment Algebra ($\mathbb{C}^4$)

Before a commitment happens, the state of the system is described by something with four complex-valued components. This is the pre-commitment state space — mathematically, $\mathbb{C}^4$, the four-dimensional complex vector space — with evolution given by unitary transformations (U(4)).

This is where quantum mechanics comes from, in VERSF. The four-dimensional structure comes from the two independent binary degrees of freedom on the fold ($2 \times 2 = 4$ states). Unitary evolution comes from reversibility of the pre-commitment dynamics. And the complex numbers are not imposed — they are the unique minimally admissible scalar field. Specifically:

Real numbers fail because the real orthogonal group does not support a global U(1) phase commuting with all admissible observables. Real-linear evolution admits continuous rotations but cannot support independently accessible relative phases on composite states; the resulting interference statistics collapse into classical statistical mixtures at the level of admissible facts. This is the Stueckelberg–Araki result, derived here rather than postulated.

Quaternions fail because quaternionic phases, acting within the four-state interface structure, either produce the same fact outcomes as the complex representation (making them observationally redundant) or require degrees of freedom beyond the four-state interface (which violates minimality — the four states are the complete set of independent binary degrees of freedom).

Octonions fail because they are non-associative. Composing reversible transformations consistently — $A \circ (B \circ C) = (A \circ B) \circ C$ — requires associativity, which octonions lack.

By Hurwitz's theorem, $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, and $\mathbb{O}$ are the only normed division algebras over the reals. Three of the four fail. Only $\mathbb{C}$ remains.

From $\mathbb{C}^4$ to the Hilbert Space

Naming $\mathbb{C}$ as the scalar field and the state space as four-dimensional is not yet a full Hilbert-space structure. A Hilbert space requires an inner product — a specific sesquilinear form $\langle \cdot, \cdot \rangle : \mathbb{C}^4 \times \mathbb{C}^4 \to \mathbb{C}$ under which states acquire norms and probabilities. The question is whether the admissibility conditions of Part 1 force that inner product, or whether it is an additional input.

The chain is the following. (i) The fold's four binary states $(\sigma, \omega) \in \{0, 1\} \times \{-1, +1\}$ identify a four-dimensional basis, but as labels they have no norm or angle relations. (ii) Pre-commitment dynamics are reversible, which means they preserve some structure on the state space — something has to be conserved under the unitary evolution, or "unitary" has no content. (iii) Admissibility (Part 1, commitment 4) requires that any two inner products producing the same fact distributions across all admissible protocols are identified — there is no physical content in distinctions between inner products that observationally agree. (iv) The combination of reversibility plus admissibility narrows the inner product down to the standard positive-definite sesquilinear form on $\mathbb{C}^4$ up to an overall normalisation, which is itself fixed by requiring the Born-rule probabilities to sum to 1.

Specifically: reversibility forces conservation of *some* positive-definite quadratic form (otherwise generic evolutions would blow up or decay, contradicting reversibility); $\mathbb{C}$-linearity plus Hermiticity forces that form to be sesquilinear; and admissibility forces all admissibly equivalent forms to be identified, leaving a unique canonical representative. The standard inner product on $\mathbb{C}^4$ is that representative. Alternatives either break reversibility (non-positive-definite forms), break $\mathbb{C}$-linearity (bilinear rather than sesquilinear forms), or fail to produce distinct fact distributions from the standard form (in which case they are admissibly equivalent, not genuine alternatives).

This argument is structural rather than fully formal. The current papers establish $\mathbb{C}^4$ with unitary evolution rigorously; the specific derivation that forces the inner product from admissibility alone, without any additional structural input, is listed in Part 6 as an in-progress item. Readers should understand the Hilbert space structure of the pre-commitment sector as *strongly motivated but not yet standalone-derived*, in the same sense that $\eta = 1$ is strongly motivated in the commitment-barrier derivation but not yet given a standalone proof.

What follows downstream from the Hilbert-space structure — superposition, interference, measurement statistics, the Born rule, the composability of systems via tensor products — does not depend on closing this specific gap. Each of those downstream features can be derived from weaker starting points. But a reader who wants to know whether the guide's "only $\mathbb{C}$ survives" statement completes the derivation should read it with the correct scope: it establishes the *scalar field* uniquely, not the *inner product* uniquely. Both are needed for Hilbert space; the scalar field is derived, the inner product is strongly constrained but still awaits a single-step derivation.

Papers: *Complex Hilbert Space from Distinguishability*, *Why a Fact-Producing Universe Must Satisfy Interference, Isotropy, and Representational Invariance*, *Quantum Mechanics as Admissibility Fixed Point*, *Measurement as Commitment*.

Layer 4: Commitment Events and the Record Field ρ(x,t)

When a commitment happens, a fact is produced. The accumulated density of committed facts across space and (emergent) time is described by the scalar field $\rho(x,t)$, called the committed record density. This is a bulk field — it lives throughout the emergent volume of space, not on the fold boundary itself.

Every observable quantity in VERSF is ultimately a functional of $\rho(x,t)$ and its derivatives. This is the Single-Source Theorem — a strong claim that deserves a sketch of how it actually works, since "all of physics from one scalar field" can sound unifyingly miraculous rather than structurally grounded.

The mechanism is this: ρ(x,t) is not just a number at each point — it is a density with gradients, time derivatives, and higher-order structure, all of which encode distinct physical information. Specifically:

Geometry emerges from spatial gradients and curvature of ρ. Variations in commitment density across space define effective distances and angles. The emergent metric is constructed from ρ and its spatial derivatives. Gravity, in VERSF, arises from gradients in ρ producing the same phenomenological effects as spacetime curvature in general relativity.

Dynamics emerges from time evolution of ρ. The rate of change ∂ρ/∂t encodes the flow of commitment events. Equations of motion for observable fields are derived from how ρ evolves under the constraints of commitment-capacity conservation (the BCB principle).

Fields emerge from specific structural features of ρ. The κ-field — the scalar field governing commitment-capacity dynamics — is a specific functional of ρ. Gauge fields emerge from closure relations in the K=7 cell structure: closed paths traversed by ρ through the cell generate holonomies, and the holonomy group reproduces the gauge content of the Standard Model (derivation developed in *The Standard Model from Hexagonal Geometry*). Particle content emerges from stable localised solutions of the equations governing ρ.

So the Single-Source Theorem is not "everything mysteriously unifies" — it is "the single quantity ρ(x,t) carries enough independent structure (value, gradient, time derivative, higher derivatives, topological features) to generate distinct classes of observable physics." The unification is structural, not miraculous.

A further question a careful reader will press: even granting that ρ(x,t) is *sufficient* to encode observable physics, how do we know it is *exhaustive*? Could there not be additional hidden degrees of freedom — fields, variables, modes — that contribute to physics without being functionals of ρ? The Single-Source Theorem takes the exhaustiveness seriously and closes it via the Saturation Theorem argument, which runs as follows.

Suppose, for contradiction, that some physical degree of freedom X exists outside ρ(x,t). Two cases. Case 1: X is observable — that is, it contributes to some measurable quantity. Then X appears in the fold's observable algebra. But the observable algebra of the fold architecture is established (by the Saturation Theorem of *Observable Closure Under the VERSF Fold*) to be $M_4(\mathbb{C})$, whose content is completely captured by functionals of ρ and its derivatives. So X is already a functional of ρ, contradicting the assumption that it is outside ρ. Case 2: X is unobservable — it contributes to no measurable quantity. Then admissibility (Part 1, commitment 4) identifies any theory with X and any theory without it, since both produce the same fact distributions. X carries no physical content and is eliminated by observational equivalence. Both cases collapse. No hidden physical degree of freedom outside ρ can survive.

This is what makes the Single-Source claim a *theorem* rather than a modelling choice: it is not "ρ happens to be enough," it is "nothing else can exist that isn't already ρ." The structural core of the argument is the combination of closure (the fold exhausts the observable algebra) and admissibility (unobservable structure is physically vacuous). The Single-Source Theorem is the downstream statement that makes this exhaustiveness explicit at the level of the record field.

The gravity derivation supplies an explicit physical realisation of this abstract exhaustiveness claim. In *Gravity from Fold Density Gradients*, gravitational sourcing is derived as proportional to ρ_bound, mass is identified with committed fold density, and the resulting field equations in the weak-gradient regime depend only on ρ and its derivatives — no additional sources, no external matter terms, no auxiliary fields. What the Single-Source Theorem establishes in the abstract — that all physics is a functional of ρ — the gravity derivation realises concretely in the specific case of gravitational dynamics, which is historically the place where theories are most tempted to smuggle in auxiliary structure (stress-energy tensors built from independent matter fields, bare cosmological constants, etc.). The concrete realisation strengthens the abstract theorem from a structural claim into one whose derivational consequences are already being worked out in at least one observable sector.

A note on the relationship to closure. The Single-Source Theorem and the closure result are independent. The closure result — that any admissible theory producing stable facts must contain the fold architecture — is established from the operational axioms A1–A3 without invoking ρ(x,t) as a load-bearing primitive. The Single-Source Theorem is a separate downstream strengthening: having established the fold, and having derived the commitment event as its operational output, one can then identify ρ(x,t) as the unified carrier of all observables. A reader who accepts closure but wants to withhold judgement on the Single-Source identification until further work is free to do so; neither result depends on the other.

Papers: *The Theory of Fact Production*, *Deriving the Commitment Capacity Density*, *Single-Source Theorem*.

Layer 5: Emergent Time, Space, and Causality

Time emerges from the ordered sequence of commitment events. Before facts exist, there is no physical time. Once facts start happening in ordered sequences, the ordering is what time *is*.

This ordering is partial and observer-dependent, not absolute. Different observers in different regions of emergent spacetime may disagree on the ordering of spacelike-separated commitment events — and this is exactly what special relativity requires. VERSF does not predict a universal absolute time; it predicts the emergence of a causal ordering that is locally well-defined but globally only partial, recovering Lorentz invariance at scales large compared to the commitment scale. The separate papers on proto-time and emergent Lorentz invariance establish this explicitly.

Space emerges from the geometric relationships between commitment events. The effective manifold structure we recognise as three-dimensional space is built up from the commitment substrate at scales much larger than the fundamental commitment scale.

Causality emerges from the directed relationships between irreversible facts. Fact A causally precedes fact B when A's occurrence is necessary for B's commitment event. Time and causality are two aspects of the same structure — the ordering of irreversible commitments.

The Emergence of Physical Time

The opening claim of this layer — that time emerges from the ordered sequence of commitment events — compresses a two-stage structural derivation that deserves fuller treatment. A reader encountering VERSF through the depth paper *Depth as a Derivative of Time* may leave with the impression that temporal sequencing is simply taken as primitive in the framework, because that paper explicitly defends sequencing as an "acceptable primitive" on epistemic-irreducibility grounds (§5.3 of that paper). VERSF does not stop there. Within the broader programme, sequencing itself is derived, and the derivation is one of the framework's load-bearing structural results.

Stage 1: Proto-time is reversible and pre-physical. Below the fold's commitment interface (Layer 2), the pre-commitment sector supports reversible unitary evolution on $\mathbb{C}^4$. This evolution has an ordering — *proto-time* — but it is not physical time in the sense physics measures. Proto-time carries no arrow (reversible dynamics do not single out a direction), no duration (there is no clock yet to measure against), and no metric structure (durations are not yet well-defined). Proto-time is what the framework has when it has dynamics but not yet facts. The detailed treatment is in Layer 1 and in the papers on proto-time.

Stage 2: Commitment converts proto-time into physical time, with arrow. A commitment event — an irreversible transition from reversible possibility to committed fact — is what converts proto-time into physical time. The arrow of time is not a statistical coincidence that could in principle have gone the other way; it is the direction of commitment itself. Each commitment event externalises its discarded alternatives to the Void (Layer 0), and that externalisation is one-way: the Void absorbs, but it does not return. The structural consequence is that the sequence of commitment events is ordered in a direction that cannot be reversed without undoing what the Void has absorbed — which, by the Void's capacity structure, it cannot do. Physical time is the ordering of this commitment sequence, and its arrow is inherited directly from the irreversibility of commitment, not from initial-condition assumptions, Past Hypotheses, or coarse-graining asymmetries. The arrow is architectural, not accidental.

Why this matters for the rest of the layer. When VERSF talks about "temporal sequencing" in the depth-reconstruction material below — and wherever else the programme invokes sequencing — the sequencing in question is *physical time*, produced by commitment. It is not a raw primitive. This places VERSF's account one level deeper than frameworks that accept sequencing as a foundational posit. VERSF accepts the sequencing those frameworks use, but locates its origin in the commitment events that generate the ordering and, through their irreversibility, give it direction. Depth-as-derivative-of-time in VERSF is therefore depth-as-derivative-of-derived-time — a two-stage reconstruction from the commitment substrate, not a one-stage reconstruction from sequencing-as-primitive.

Observer-dependence and the recovery of Lorentz invariance. The ordering of commitment events is partial, not total. Two spacelike-separated commitments do not have a frame-independent ordering — different observers sample the commitment sequence differently, producing locally well-defined partial orderings that agree on causally connected pairs and disagree on causally independent ones. At scales large compared to the commitment scale ξ, this partial-order structure is what recovers the relativity of simultaneity of special relativity. Lorentz invariance is emergent, not imposed: it appears when the isotropy of the Void's response to commitment dynamics is combined with the invariant propagation speed c derived later in this layer. The separate papers on emergent Lorentz invariance and the proto-time / physical-time distinction develop this explicitly.

The full dependency chain. Putting the stages together, time in VERSF sits at the end of a dependency chain: proto-time (reversible, pre-physical) → commitment event (irreversible transition) → physical time (ordered, directed) → partial causal ordering → invariant propagation speed c → Lorentzian structure at macroscopic scales. Each step rests on the one before; none is imposed from outside. What finite observers call "time" is the end of this chain, not its beginning. Layer 1 gives proto-time in detail, Layer 2 gives the fold and the commitment mechanism, and the subsections below give the causal and metric structure that sits above physical time.

From Discrete Commitments to Continuous Geometry

A sharp question arises at this point. If spacetime is emergent and the substrate is discrete, how does a smooth differentiable manifold arise at all? The answer is coarse graining — but coarse graining that is defined causally rather than spatially, since space itself is what is being derived. Averaging over pre-existing spatial volumes would be circular.

At the substrate level, the framework contains individual commitment events and the causal relations between them. A commitment event's bounded causal neighbourhood is the primitive structure that exists before geometry. The programme identifies a fundamental coarse-graining scale R* — the minimum size of causal neighbourhood over which commitment-event statistics become well-defined. At scales ≥ R*, statistical aggregation over causally connected commitment events produces a smooth entropy field s(x), by the law of large numbers.

Geometry emerges from s(x) in three distinct steps that should not be collapsed:

Step 1: Coarse graining defines the entropy field. Aggregation over causal commitment neighbourhoods at scale R* produces the smooth field s(x).

Step 2: The TPB law drives dynamics. The commitment-flux current is governed by gradients in s(x). This follows from Bit Conservation and Balance combined with the capacity constraint; it is not posited.

Step 3: The Void responds to entropy gradients under finite-capacity constraint. What we recognise as distances, durations, and curvature is the Void's structural response to gradients in s(x), not variation in the field itself. The entropy field drives; the Void deforms; the deformation is bounded by the substrate's finite capacity. In the continuum limit ($L \gg R^*$) this response defines an effective pseudo-Riemannian structure whose detailed construction — the explicit form of the metric tensor, the conditions for Lorentzian signature, and the weak-gradient reduction to Newtonian gravity — is developed in *Record-Theoretic Emergence of Spacetime Geometry and Gravitational Dynamics* and *Where the Speed of Light Comes From*.

Depth as Temporal Reconstruction

A structural clarification is needed about the origin of spatial depth in this framework. The positive account is developed in *Depth as a Derivative of Time*, and the guide locates it here because it is the natural companion to *Depth Is Not a Direction*: the first paper establishes what depth is *not* (a primitive spatial direction, a parameter of descriptive scale), the second establishes what depth positively *is* (a coordinate reconstructed from temporally ordered update dynamics).

No single configuration Σ_n contains a depth coordinate. Depth is not a primitive geometric direction; it is a derived quantity constructed from temporal sequencing. Specifically, depth is a coordinate constructed from structured differences between successive configurations under the update operator $\mathcal{T}_n$. The ordering of commitment events supplies the sequencing; the accumulation of distinguishability changes across that sequencing defines how far two configurations are from coarse-graining to the same operational state. Formally, this reconstruction is captured by a depth functional $\mathscr{F}$ defined over trajectories of the update sequence, with the depth between two initial states given by the minimum update count at which they become operationally indistinguishable under cumulative update ($Z = \kappa \cdot z$ in the notation of the depth paper). Observable spatial relations correspond to equivalence classes of trajectories under this reconstruction, ensuring that the derived depth coordinate has operational meaning for finite observers: what an observer measures as spatial separation is an invariant of the equivalence class, not of any particular trajectory within it.

The relation is analogous to motion in film: no individual frame contains motion, yet motion emerges from structured differences across frames. In the same way, no single configuration contains depth, yet depth emerges from structured differences across updates. The analogy is not loose — in both cases, what appears to be a primitive feature of the thing being observed (motion, depth) is in fact a derived feature of its temporal structure, and the derivation is mathematically precise rather than merely suggestive.

The implication for the rest of Layer 5 is direct. Spatial depth is not introduced as an additional dimension; it emerges as a coordinate encoding inter-slice distinguishability structure. Geometry is the stable encoding of structured change under finite-capacity constraint — built not from static structure but from the way successive configurations differ and the constraints those differences respect. The metric schematic written down in the next subsection, and the invariant speed derived in the one after, both operate on this reconstructed spatial backbone rather than on a pre-given three-dimensional manifold.

A note on where VERSF goes deeper than the source paper. *Depth as a Derivative of Time* takes temporal sequencing as primitive within its own axiom set (its §5.3 defends this explicitly on operational grounds). VERSF carries the richer account developed in the *Emergence of Physical Time* subsection above: physical time is itself derived from commitment, with its arrow inherited from commitment irreversibility, not accepted as a foundational posit. The depth-reconstruction argument is unchanged — the ordering of commitment events is exactly the sequencing that the depth functional $\mathscr{F}$ operates on — but in VERSF the dependency chain runs one step deeper: commitment → physical time → depth reconstruction. Depth-as-derivative-of-time and time-as-derivative-of-commitment are consecutive stages of a single structural argument that the guide carries in full, even where individual papers carry only one stage of it.

The full structural dependency. Stated in one chain, the dependency runs:

commitment → temporal ordering → inter-slice differences → correlation structure → reconstructed depth → emergent geometry

Each stage is not independent but derived from the previous. Geometry is therefore not constructed directly from matter or fields, and space is not a pre-given arena within which dynamics unfold. Both are products of structured distinguishability accumulated across temporally ordered configurations, under the finite-capacity constraint of the Void. This is the ontological backbone of Layer 5, and every downstream result in the layer — the metric schematic, the invariant speed, the status of gravity, the Structural Clarifications — is a consequence or refinement of this single chain.

The Schematic Form of the Emergent Metric

A reader trained in general relativity will reasonably want to see the emergent metric written down, at least schematically, rather than only characterised verbally. The following schematic captures the structure without committing to the specific coefficients, which are fixed in the papers cited.

At scales $L \gg R^*$, the effective metric $g_{\mu\nu}(x)$ takes the form:

$$g_{\mu\nu}(x) = A(s)\, \eta_{\mu\nu} + B(s)\, \partial_\mu s\, \partial_\nu s + \text{higher-derivative corrections}$$

Here $\eta_{\mu\nu}$ is the local Minkowski reference form (set by the causal structure of commitment events: the time-direction inherits its asymmetry from the irreversibility arrow of commitment, the space-directions from the symmetric exploration of the pre-commitment sector). $A(s)$ is a conformal factor encoding how proper intervals stretch or compress under local entropy loading — regions of high s have suppressed commitment-flux rates, which finite observers register as temporal dilation and spatial compression. $B(s)$ is an anisotropic response coefficient encoding how the Void deforms asymmetrically along the gradient direction of s — commitment flux is easier along $-\nabla s$ (down-gradient) than along $+\nabla s$ (up-gradient), and this asymmetry is what produces the gravitational attraction of mass distributions (high ρ, high s) toward one another.

Three features of this schematic are worth stating explicitly.

First, $g_{\mu\nu}$ is a *response field*, not the smoothed field. The smoothed field is s(x); the metric is the Void's structural response to gradients in s. Conflating the two is the most common misreading of VERSF's emergence story. In standard induced-gravity frameworks (Sakharov, Verlinde, and others in the entropic-gravity tradition) the metric is treated as fundamental and entropy is derived from it; VERSF runs the arrow the other way. s(x) is fundamental (it is coarse-grained commitment statistics); $g_{\mu\nu}$ is derived from the Void's response.

Second, the coefficients A(s) and B(s) are structurally constrained rather than free, but their explicit functional forms are not yet fully derived in the programme. What is established is that they are functionals of the Void's constitutive response to entropy loading, and that response is in turn determined by the $K = 7$ closure-cell structure and the TPB law. What is *not* yet established in closed form is the specific derivation of A(s) and B(s) from those structural inputs — the record-theoretic emergence paper gives partial results and the weak-gradient limit, but the full forms across all regimes remain an open technical item flagged in Part 6. Readers should understand "A and B are functionals of Void response" as locating where the derivation has to happen, not as the derivation itself.

Third, the Lorentzian signature $(-, +, +, +)$ of $g_{\mu\nu}$ is not imposed. It emerges from the asymmetry between temporal accumulation (irreversible, directed — inherited from the commitment arrow at the fold level) and spatial propagation (reversible, bidirectional — inherited from the pre-commitment sector's symmetric exploration). The signature is a consequence of the fold's internal structure, not an external choice.

Fourth — and this completes the connection between geometry and dynamics — the causal structure of $g_{\mu\nu}$ is set by the invariant propagation speed derived in the next subsection. The metric does not encode geometry alone; it simultaneously encodes the maximal rate at which distinguishability can propagate through the substrate, with that rate fixed structurally rather than postulated. Geometry and causal structure in VERSF are two aspects of the same Void response, not independent layers.

In the weak-gradient limit ($B(s)|\nabla s|^2 \ll A(s)$), this schematic reduces to an effective Poisson-like equation for the gravitational potential, recovering Newtonian gravity. In the stronger-gradient regime, curvature terms appear that track general-relativistic behaviour in the regimes tested to date (including $c_T = 1$, which is a theorem rather than an assumption in VERSF). The full non-linear reduction to the Einstein equations — specifically, an explicit expression of the form $G_{\mu\nu} = f(\rho, \partial\rho, \partial^2\rho, ...)$ together with a stress-energy relation and the intermediate linearised-GR structure — is still in progress and is flagged in Part 6 as a principal open technical objective.

What the schematic therefore establishes, and what it does not, should be stated clearly. It establishes that the emergence is of the right structural *kind*: an effective pseudo-Riemannian metric built from the coarse-grained entropy field and its gradients, with Lorentzian signature inherited from the commitment structure, reducing correctly in the weak-gradient limit. It does not yet establish the full derivation. The functional forms of A(s) and B(s), the complete set of higher-derivative corrections, the stress-energy relation, and the step-by-step reduction to the linearised and then full Einstein equations are open. A GR-trained reader should read the schematic as "this is the right form for what must be derived," not as "this is the derivation."

Readers who want the specific weak-gradient results and the directions the full derivation is being pursued should read the record-theoretic emergence paper directly.

The Invariant Speed as a Structural Constraint

The existence of a finite, frame-independent maximum speed of causal propagation — the speed of light c — is often introduced in textbooks as a postulate. Within VERSF it is derived, as a structural consequence of three constraints on the fact-producing substrate: finite distinguishability (records require nonzero resolution capacity; A1 in the admissibility axioms), bounded local update (causal influence propagates through adjacent substrate elements only; a consequence of fact-production being a local commitment event), and causal consistency (record ordering must be globally non-contradictory; no fact can be part of a cycle that contradicts it). Together these imply that causal influence can propagate at most one adjacency step per fundamental update. In the continuum limit this defines a maximum propagation speed:

c = (minimal spatial distinction) / (minimal temporal distinction)

This is not an additional postulate but a necessary consequence of requiring stable, causally consistent fact formation within a finite-capacity substrate. If such a propagation bound did not exist, causal ordering would become inconsistent and stable fact formation would fail. The existence of a finite invariant speed is therefore structurally required. The invariance of this propagation speed across all admissible observers forces the Lorentzian structure of the emergent spacetime: causality constrains symmetry, and symmetry constrains geometry. The causal cone structure from which Lorentz symmetry emerges at the macroscopic limit follows from this speed, and the emergent metric's causal structure (flagged in the "Fourth" observation above) is fixed by it. A fuller derivation is given in *Where the Speed of Light Comes From*; the point here is that readers should not import the assumption of a postulated c when reading about emergent geometry — the invariant speed is internal to the same structural argument that produces the metric.

Depth Is Not a Spatial Direction

Coarse-graining is a central operation in this layer: the record field ρ is coarse-grained at the scale R^* to produce the entropy field $s(x)$. A reader familiar with renormalisation group methods, holographic duality, or tensor-network constructions may be tempted to treat the coarse-graining scale itself as an additional spatial direction — a "depth" into which scale-dependent physics unfolds, alongside the three familiar spatial directions. VERSF rejects this move on structural grounds, and the rejection matters for reading the rest of the framework correctly.

The four structural criteria that a parameter must satisfy to qualify as a physical spatial direction are: (C1) an intrinsic metric, (C2) locality, (C3) propagation through intermediate positions, and (C4) reversibility. Renormalisation depth fails all four. It has no scheme-independent metric — the "distance" between two resolutions depends on the choice of coarse-graining procedure. It has no locality in the required sense — "degrees of freedom at depth z" are not independent of "degrees of freedom at depth 0"; they are derived from them. It supports no propagation — no object traverses intermediate depths as a trajectory, because coarse-graining is not a dynamical evolution but a re-description of the same underlying state. And it is constitutively irreversible — non-injective by definition, because distinct fine-grained states can coarse-grain to the same effective description. The full argument is given in *Depth Is Not a Direction*, including a no-go theorem establishing that non-injective coarse-graining semigroups cannot define spatial directions.

Readers familiar with the positive account developed above may ask whether this no-go renders all depth coordinates unphysical. It does not. Unlike renormalisation depth, which indexes descriptive scale, the reconstructed depth coordinate of *Depth as a Derivative of Time* arises from physically realised update dynamics. It is therefore not a bookkeeping parameter but an operationally defined quantity derived from inter-slice evolution, and it is associated with an intrinsic metric (the physical depth $Z = \kappa \cdot z$), with locality (the propagation bound κ is local on each slice), with genuine propagation (through the sequence of updates rather than through descriptive rescaling), and with a well-defined sense of reversibility at the level of physical dynamics even when effective coarse-graining is non-injective. The two should not be conflated. The coarse-graining that fails as a spatial direction is descriptive; the sequencing that succeeds as a spatial direction is physically realised.

For VERSF specifically, this has three consequences. First, the scale R* introduced in this layer is a statistical resolution scale, not a spatial direction. It is the scale at which coarse-graining produces a smooth field; it does not extend space into a fourth dimension. Second, where VERSF derives three-dimensional space as an emergent geometric structure, that derivation proceeds from correlation structure in the record field — from the patterns of how committed facts relate to one another across the substrate — where those correlations themselves arise from structured differences across temporally ordered configurations, in the manner developed in *Depth as a Derivative of Time* and summarised in the *Depth as Temporal Reconstruction* subsection above. The "third dimension" in VERSF is reconstructed from correlation structure supplied by temporal sequencing, not imported from the scale parameter and not stacked onto a two-dimensional interface as a primitive direction. Third, VERSF's position on holographic bulk-reconstruction constructions is specifically the one developed in *Depth Is Not a Direction*: the geometric interpretation of the radial coordinate in AdS/CFT and MERA-like constructions is emergent and regime-dependent, arising from the organisation of boundary information rather than from a primitive spatial direction. Bulk geometry is a representation of information structure; it is not an additional arena in which independent physics happens.

This is not a refusal of holographic and tensor-network techniques — they are powerful and indispensable representations. It is a precise statement of what they represent: scale-organised correlation structure of boundary data, which admits a geometric interpretation in specific regimes, not an ontologically primitive spatial dimension.

Two scales matter, and they are not identical. R* is the minimal statistical neighbourhood at which aggregation produces a smooth field. ξ is the minimal commitment-capable domain at which coherent commitment dynamics actually occur; it is also the scale at which the quartic capacity threshold $\chi(L) = \rho L^4/\hbar c$ becomes operationally relevant. R* is a scale of statistics; ξ is a scale of physics. Both sit well above the Planck scale, and their ratio is determined by the K=7 cell structure.

Below R*, no geometric description is admissible — only discrete commitment structure and causal relations exist. Above it, smooth spacetime emerges as the large-scale statistical limit of a finite-resolution, fact-producing substrate, expressed through the Void's response to entropy gradients. The coarse-graining technique is standard — it mirrors the hydrodynamic limit of lattice systems and effective-field-theory logic from condensed matter. What is distinctive to VERSF is that coarse graining is defined causally rather than spatially, the smoothed field is an entropy density whose dynamics follow TPB, and geometry is the Void's response rather than the smoothed field itself.

Current Status of Gravity in VERSF

A compact summary of where the gravity derivation stands, for readers who want to assess the current state without tracking through the individual papers. *Field law:* the Poisson equation for the gravitational potential in the weak-gradient limit is derived as the unique admissible form compatible with admissibility constraints, within the class of local, isotropic, linear, second-order, scale-free field equations. *Inverse-square law:* derived as the unique admissible vacuum solution within the same class. *Gravitational coupling:* the constant G is reduced to a single structural normalisation $G = (\lambda/\mathcal{C})\cdot\xi^2 c^3/\hbar$, in which ξ is the derived coherence scale, $\mathcal{C}$ is the capacity coefficient fixed by the K = 7 closure structure, and λ is a residual efficiency factor encoding how efficiently committed distinguishability propagates through the fold closure architecture. *Weak-field gravity:* fully recovered in the Newtonian limit. *Tensor structure:* the structural argument for how the scalar potential is promoted to a rank-2 symmetric field is in place; the detailed derivation is in progress. *Full General Relativity:* not yet derived; the explicit reduction from the entropy-field dynamics to the linearised and then full Einstein equations, together with the stress-energy relation, is a principal open technical task flagged in Part 6. What the programme has achieved in gravity is substantial but not yet complete: the unique admissible form of gravitational sourcing is fixed, the coupling is reduced to a single one-parameter open closure, and the weak-field limit is recovered; what remains is the step-by-step upgrade from schematic reproduction to full derivation.

Structural Clarifications

Five results developed in this layer are worth isolating, because they are the ones most commonly misread and because stating them together makes their combined force visible.

- The invariant speed c is structurally required, not postulated. It is a necessary consequence of stable, causally consistent fact formation within a finite-capacity substrate.
- Renormalisation depth is not a spatial dimension. It fails all four structural criteria for spatiality (intrinsic metric, locality, propagation, reversibility) and is a descriptive resolution parameter rather than a coordinate.
- The third spatial dimension is emergent from correlation structure, not stacked onto a two-dimensional interface as an additional primitive direction.
- Geometry reflects constraints on distinguishability, not a pre-existing manifold. The metric is a response field; the manifold is what emerges when that response is coarse-grained.

- Physical time is not primitive but derived from commitment. Its arrow is inherited from commitment irreversibility, not from initial-condition asymmetries or statistical accidents. Where other frameworks take temporal sequencing as a foundational posit, VERSF locates its origin one level deeper, in the commitment events that produce the ordering.

Taken together, these five points pin the ontology of this layer: a 2D fact-producing substrate, with physical time as the ordering of commitment events and space as the three-dimensional correlation structure emergent above it, all organised under the structural limit c and the finite-capacity constraint of the Void.

In a single sentence: **time provides ordering, differences generate structure, and geometry is the stable encoding of that structure.**

Papers: *Where the Speed of Light Comes From*, *Record-Theoretic Emergence of Spacetime Geometry and Gravitational Dynamics*, *Gravity from Fold Density Gradients*, *Why Time Builds Space (Not the Other Way Around)*, *Causation Is Not a Wave*, *A Hidden "Middle Scale" in the Universe — and Why It Matters* (the mesoscopic scale ξ), *The Planck Time as a Certification Floor for Emergent Temporal Metrics*, *Two Descriptions of Reality — The Coherence Scale as a Commitment Threshold*, *Depth Is Not a Direction*, *Depth as a Derivative of Time*.

Layer 6: K=7 Substrate Architecture

Once you have the fold as a two-dimensional commitment interface, the next question is how many of them fit together to form the local substrate. VERSF claims the answer is exactly seven — K=7 — forced by a combination of simplicial admissibility constraints and the algebraic properties of the Fano plane (the smallest finite projective plane, whose symmetry group is GL(3,2) with 168 elements).

This is the layer at which the framework becomes quantitatively predictive. It is also the layer where critics are most likely to suspect tuning — so the question "why not 6, or 8?" deserves a preview answer rather than a reference.

The short version: simplicial admissibility imposes specific closure conditions on how interfaces can be joined. For $K \leq 6$, the closure structure has too few independent constraints — the resulting substrate admits degenerate configurations that do not produce stable fact formation. For $K \geq 8$, the closure structure is overdetermined — no consistent simplicial configuration exists that satisfies all the constraints simultaneously. $K = 7$ is the unique value where the constraint system is exactly saturated: neither underdetermined nor overdetermined. This is tied to the Fano plane (7 points, 7 lines, each pair of points on a unique line), whose automorphism group GL(3,2) provides the symmetry structure of the K=7 cell.

The No-Go Theorem paper establishes this through three independent routes, which provide the multiple-route overdetermination that distinguishes a derived value from a tuned one. Constraint counting: the closure system of interface junctions yields a specific number of independent relations versus free parameters, and K=7 is the unique value where these balance. Group-theoretic saturation: only GL(3,2) provides the automorphism structure required for a simplicial substrate supporting fact-producing commitment interfaces, and GL(3,2) is specifically the Fano-plane symmetry group with K=7. Simplicial dimension: the minimal simplicial complex supporting codimension-1 commitment interfaces compatible with the other substrate constraints has exactly seven cells. A single argument for K=7 could be coincidence; three independent arguments landing on the same number is the signature of genuine structural necessity.

The K = 7 No-Go Theorem (unified statement)

The three routes can be stated as a single theorem, which is the form the No-Go paper takes at its core. This consolidates what "K = 7 is forced" actually means as a mathematical claim.

Theorem (K = 7 No-Go). Let S be a simplicial relational substrate satisfying:

- (T1) *Non-degenerate closure.* The closure operations on S are non-degenerate: distinct interface configurations do not collapse to the same closure under composition.
- (T2) *Fact-supporting topology.* S supports codimension-1 commitment interfaces with the binary geometric data (σ, ω) derived in Layer 2.
- (T3) *Automorphism saturation.* The automorphism group of S acts transitively on closure paths, with the minimum group structure required for consistent composition of commitment events.
- (T4) *Minimal architecture.* S has the smallest cell count consistent with (T1)–(T3).

Then |S| = 7, and the automorphism group is GL(3,2) ≅ PSL(2, 7), the symmetry group of the Fano plane.

Three observations. First, the theorem is a single conjunction of four conditions forcing a single number; it is not a conjunction of three independent numerical arguments that happen to agree. The three "routes" in the earlier paragraph are three different proofs that the four-condition system has |S| = 7 as its unique solution: constraint counting establishes it from (T1) and (T4), group-theoretic saturation from (T3), and simplicial dimension from (T2) plus (T4). The multiple-route character is what makes the result robust; the theorem itself is single-claim.

Second, dropping any of (T1)–(T4) changes the result. Without (T1) the system admits degenerate K < 7 configurations. Without (T2) the substrate can have fewer cells if it does not support fact-producing interfaces. Without (T3) the automorphism structure is too weak to support consistent composition. Without (T4) K ≥ 7 are all admissible and no unique value is forced (higher-K realisations reduce to K = 7 at the fact-supporting level, but (T4) is what pins the minimal case). Each condition does specific work.

Third, the theorem is a statement about the *minimal fact-supporting substrate*, not about substrate dimension in an absolute sense. A universe can have additional structure beyond the K = 7 core — what the theorem establishes is that any fact-supporting core it contains must be K = 7 at the structural minimum. This is the same strictness asymmetry that runs through the programme: downstream structures can add further specificity, but the core they rest on is pinned.

This is what distinguishes K = 7 from a numerological fit. A fit says "the number 7 is what matches the data." The theorem says "anything less than 7 fails specified structural conditions, anything more than 7 reduces to 7 at the structural core, and 7 itself is pinned by the intersection." The data never enters. Whether the downstream quantitative consequences (α, Λ, masses) match observation is an empirical test of the theorem's consequences, not of the theorem itself.

The Convergence Theorem: Six Labelled Constraints, Four Mathematical Structures

The No-Go theorem above establishes K = 7 from a specific set of structural conditions on the minimal fact-supporting substrate. A separate and complementary result — the Convergence Theorem developed in *K = 7: The Constraint Dimensionality of Stable Physical Reality* — establishes K = 7 from a broader set of independently-motivated constraints drawn from six different physical and mathematical domains. Where the No-Go theorem asks "what is the minimal substrate?", the Convergence Theorem asks "what integer satisfies all these independent constraints simultaneously?" — and answers with the same K = 7.

The six labelled constraints are: (C1) geometric closure, from the Honeycomb Theorem and a rank-nullity argument on the constraint matrix of the hexagonal interface; (C2) boundary entropy consistency, from the Bekenstein–Hawking requirement that a causal 2D boundary encode the 66 = C(12,2) pairwise relations among the 12 Standard Model gauge generators; (C3) minimal gauge encoding, from the minimality condition $2^K \geq C(12,2) = 66$; (C4) topological integrality, from Chern–Weil integrality of the electromagnetic winding number on the de Sitter horizon; (C5) combinatorial balance, from the Diophantine equation arising in the BCB proton-mass derivation; (C6) axiomatic admissibility, from PAR (Pre-Factual Algebraic Reversibility) plus gauge-restricted Compositional Completeness.

An honest independence analysis establishes that these six labels reduce to four genuinely distinct mathematical structures. C2, C3, and C6 share the same core inequality $2^K \geq C(12,2) = 66$, approached from three different physical framings (thermodynamic, spinorial, axiomatic); they should be counted as one constraint with three independent motivations, not as three independent constraints. C1 and C6 share the hexagonal-tiling conclusion, reached by independent means. C5 shares no assumptions with any other constraint. The four distinct structures are therefore: geometric closure (C1, C6 geometric arm), binary gauge encoding (C2, C3, C6 encoding arm), topological integrality (C4), and combinatorial balance (C5). Four genuinely distinct mathematical frameworks converging on a single integer is a non-trivial result — weaker than "six independent routes," stronger than "one argument dressed six ways."

Among these, the cleanest single result is the Diophantine equation of C5:

$$2^K + 8 = N(N - 1) / 2$$

This equation has a unique positive-integer solution: K = 7, N = 17. The left side grows exponentially in K; the right side is quadratic in N; exhaustive checking through K = 30 confirms uniqueness, and an asymptotic argument rules out further solutions. For K = 7, the right side gives 17 × 16 / 2 = 136 = 128 + 8 = $2^7 + 8$, which verifies. The offset +8 is not fitted: it is derived from the specific BCB coupling structure between the electron block and the baryon core (three inter-subgroup pairings contributing 2 states each under the Z_2 balance condition, plus two inter-block closure modes — total 8). Changing the offset to +7 or +9 produces equations with no integer solutions. C5 is therefore fully unconditional and integer-exact: no shared assumptions with any other constraint, no numerical approximation, no free parameters. It is the single strongest result in the K = 7 programme.

Honest Status of the Six Constraints

Different constraints sit at different levels of completeness, and a reader evaluating K = 7 should know which is which. A concise summary:

- **C5 (combinatorial balance).** Fully unconditional. The Diophantine equation and its unique solution are integer-exact and independently verifiable.
- **C3 (minimal gauge encoding).** Rigorous as a minimality condition, conditional on N_gen = 12 (an empirical input from the Standard Model). If N_gen is derived from the VERSF axioms rather than imported, C3 becomes fully unconditional.
- **C2 (boundary entropy consistency).** Lower bound derived; upper bound requires explicit calculation of the independent holonomy degrees of freedom on the simplicial boundary complex (an open task).
- **C6 (axiomatic admissibility).** Lower bound derived from completeness; upper bound is a structural argument pending explicit representation-theoretic calculation on the hexagonal interface symmetry group.
- **C1 (geometric closure).** Established. The rank-6 / nullity-1 structure of the hexagonal cell's response matrix is proved in *Completing the Interface Bridge* (Appendix A) via the graph-Laplacian structure of the hub-and-boundary graph G_7, together with the Honeycomb Theorem for tiling selection and an orbit-stabiliser enumeration for the $K \leq 7$ bound.
- **C4 (topological integrality).** Conditional on the coherence-scale derivation $\xi(K)$ developed in the companion *Geometric Closure* paper. Its force lies in the ~53-order-of-magnitude gap between adjacent integer values of K, which makes accidental agreement between K = 7 and the observed Λ–α band effectively impossible.

C5 alone suffices to pin K = 7 as a mathematically unique integer. The other five constraints are each independently motivated, several of them conditional on companion derivations, and collectively they show that the same integer sits at the intersection of four distinct mathematical structures. This is the structural content of the K = 7 result — not "six independent proofs" but "four mathematical structures, one integer-exact unconditional constraint, and an architecture in which downstream derivations progressively tighten the conditional ones."

Downstream, K=7 determines specific numerical values: the fine-structure constant, the cosmological constant, particle masses, coupling strengths, and the observable consequences of gravity. If K were not exactly 7, these constants would shift — which is what makes K=7 empirically testable rather than merely assumed. What a reader should take from Layer 6 is that the integer K = 7 is pinned by two convergent structural results — the No-Go theorem on the minimal fact-supporting substrate, and the Convergence Theorem on six independently-motivated constraints reducing to four mathematical structures — with C5's Diophantine equation providing the single fully-unconditional anchor. Downstream quantitative consequences are developed in Layer 8, with honest status accounting of which are derived, which are partially derived, and which are open.

Papers: *A No-Go Theorem for Non-Simplicial Relational Substrates — K=7*, *K=7: The Constraint Dimensionality of Stable Physical Reality*, *K=7: The Hidden Structure Behind Physical Reality*, *Why the Standard Model Might Not Be Arbitrary After All*, *Proton Mass from BCB Constraint Dynamics*, *Geometric Closure*, *Completing the Interface Bridge*.

Layer 7: Fields, Forces, and Particles

Above the K=7 substrate sit the derived physical structures: the κ-field (a scalar field governing commitment-capacity dynamics), the gauge structure of the Standard Model (emerging from hexagonal closure geometry), the particle spectrum (emerging from fold-boundary closure modes), and the forces (four fundamental forces, each corresponding to a distinct class of closure relation).

Specific structural results at this layer include the κ-field mass coefficient $\sqrt{4/3}$, fixed by PGL(3,2) irreducibility acting on the 6-dimensional non-uniform subspace V_6 of the K=7 closure geometry. This is the kind of result that distinguishes the framework from a generic effective-field-theory treatment: the coefficient is not a free parameter fit to observation, it is forced by the representation theory of the closure-cell symmetry group acting on the relevant subspace. Similar derivations underwrite the lepton sector masses, the closure-entropy derivation of the commitment barrier Φ_c, and the structural origin of specific coupling ratios.

Papers: *The κ-Field*, *The Standard Model from Hexagonal Geometry*, *Why Physics Has Exactly Four Fundamental Forces*, *Lepton Sector Mass Derivation*, *The κ-Field Mass: Structural Derivation via PGL(3,2) Irreducibility*.

Layer 8: Predictions and Constants

At the top of the stack sit the specific numerical predictions. It is important to position these correctly: they are downstream quantitative consequences of the fixed structural framework, not part of the core uniqueness derivation. Once closure has established the fold, saturation has fixed the observable algebra, and K=7 has determined the substrate architecture, the quantitative constants follow from the geometry of the closure system rather than from new fundamental assumptions. A reader should separately assess the structural core (which either holds as a proof or fails) and the quantitative constants (which either match observation or do not).

The specific results at this layer sit at different levels of derivational completeness, and the guide states the levels honestly rather than collapsing them into a single "derived" claim.

The fine-structure constant. The bare term $\alpha^{-1}_{bare} = 2^7 \cdot 15/14 = 128 \cdot 15/14 \approx 137.143$ is derived, with no coefficient fitted to observation. The target value is $\alpha^{-1}_{obs} \approx 137.036$. The bare term exceeds the target by about 0.1, which requires a small second-order correction. The structural inputs to that correction are now established: *Completing the Interface Bridge* derives the six-channel interface structure (from the Honeycomb Theorem and the K=7 Counting Theorem), the uniform allocation $w_i = 1/6$ across channels (from the transitive $\mathbb{Z}_6$ action under isotropy), the IPR coefficient 1/6 (as a direct corollary), and the exclusion of cross-channel covariance terms from the primitive local observable algebra (by P3 locality and the link-variable structure). The second-order correction functional form is therefore $\mathcal{L}_2 = (1/6)\,\mathrm{Var}(G)$, with the 1/6 derived rather than assumed. What remains open is the phase-resolution collapse — whether the admissible set $\mathcal{C}$(K, N_loop) for the discrete phase-resolution parameter N_φ collapses to a unique function F(K, N_loop), or retains residual freedom that would need to be fixed by additional physical conditions. This is open problem O1 of *Completing the Interface Bridge* and is the specific item on which the quantitative completeness of the α-derivation now hinges. What is established: the bare term, the six-channel structure, the uniform allocation, the 1/6 IPR coefficient, the cross-term exclusion, and the non-fitted character of the formula. What is open: the collapse of $\mathcal{C}$(K, N_loop).

A note on earlier formulations. The α-derivation has passed through multiple formulations in the programme's history, and Part 6's technical exposures list one of them — the earlier framing in which $\alpha^{-1} \approx 137$ is derived to 0.08% accuracy conditional on identifying the interface-coupling scale Λ_interface with the electron mass scale m_e. The bare-term-plus-IPR framing above is the current formulation and refines that earlier result: the bare term is now unconditional (no scale match required), the correction's functional form is structurally derived from the six-channel IPR, and the live open item is O1 rather than the Λ_interface-to-m_e identification. The two formulations are successive stages of a single derivation programme, not competing derivations. Readers encountering the Λ_interface exposure in Part 6 should therefore understand it as belonging to the earlier-stage pathway; the scale-match concern does not recur in the current framing.

The cosmological constant. The functional form $\xi(K) \propto \exp(\pi \cdot 2^K / 2)$ together with $\Lambda = 3/\xi^2$ is derived, placing K = 7 uniquely in the observed cosmological band — the winding-number separation between adjacent integer values of K is roughly 53 orders of magnitude, which makes the selection of K = 7 robust against uncertainty in the proportionality constant. What is open: the proportionality constant itself, which converts the functional form into the observed magnitude, requires the coherence-scale normalisation developed in *Geometric Closure*. The structural result (K = 7 lands in the observed band) is derived; the absolute normalisation is conditional on the companion-paper derivation.

The Bekenstein–Hawking entropy. The area-linear scaling $S \propto A$ is derived, and K = 7 enters the coefficient, consistent with the BH formula's form. What is open: the exact 1/4 numerical coefficient. The VERSF discrete count gives $S \approx 7 \log 2 \cdot (A/\ell^2\_P) \approx 4.85 \cdot (A/\ell^2\_P)$, which agrees with $S\_BH = A/(4\ell^2\_P)$ in order of magnitude and in the area-linear form but differs in the overall normalisation by a factor of ~19. Closing this requires a normalisation derivation that converts between the VERSF discrete constraint count (using log 2 per bit) and the continuous entropy measure in Planck units, developed in *Completing the Interface Bridge*. What is established: area-linear scaling with K = 7. What is open: the exact 1/4.

The tensor-to-scalar ratio $r \approx 0.027$–0.033. Derived from the κ-field spectrum in *Tensor Perturbations in VERSF*. Testable by LiteBIRD around 2032, which would either confirm or rule out this specific prediction.

The commitment barrier coefficient $C^* = 3/8$. Emerges from closure geometry once K = 7 and the κ-field dynamics are fixed. Derived in *Deriving the Commitment Barrier from Closure Entropy and the Coherence Scale*.

The proton mass. The BCB derivation gives $m_p \approx 936.4$ MeV for N = 17, versus the measured 938.27 MeV (0.2% residual). The shell count N = 17 is derived from BCB closure constraints, not fitted, and is the same integer that appears in the Diophantine equation of C5. The 0.2% residual is within the accuracy expected at the current state of the derivation.

Other constants. The proton-to-electron mass ratio, specific gauge coupling ratios, and various other downstream numerical consequences are under active derivation. Each carries its own derivational-status label in the companion papers; readers should not assume that every numerical mention in Layer 8 sits at the same level of completeness as α^{-1}_{bare} or the Diophantine equation.

The overall pattern: the bare/structural terms are typically derived, the first-order approximations are typically within observational range, and the second-order corrections and exact normalisation factors are typically the open items. This is the current state of the quantitative programme — substantial progress on many fronts, honest accounting of what remains.

Papers: *The Fine-Structure Constant from VERSF Architecture*, *K = 7: The Constraint Dimensionality of Stable Physical Reality*, *From Admissibility to 1/137*, *Deriving the Cosmological Constant*, *Tensor Perturbations in VERSF*, *Where the 1/4 in Black Hole Entropy Comes From*, *Deriving the Commitment Barrier from Closure Entropy and the Coherence Scale*, *Proton Mass from BCB Constraint Dynamics*.

Part 3: The Spine of the Argument

The layer stack describes the structure. The spine is the logical chain that shows why that structure is not optional.

There are two complementary presentations of the structural argument. The first is the *admissibility chain*: a compressed, five-step account of what any theory of the observable must implement if it is to support stable, compositional records. The second is the *elimination spine*: a twelve-step derivation that takes the four starting commitments and walks them to the fold, ruling out alternatives at each step. The admissibility chain is the programme's principal structural result as presented in *The Layered Foundation of VERSF*; the elimination spine is the form the argument takes when expanded to the full fold derivation. A reader who wants the compressed structural case should take the admissibility chain; a reader who wants the detailed elimination logic should take the spine. Both arrive at the same architecture.

The Admissibility Chain (five steps)

Any theory of the observable that supports stable, compositional records must implement five structural features. Each step rules out a would-be alternative by identifying what physics loses if the step's constraint is dropped.

Step 1 — Without distinguishability, no observation. A theory positing observable content without minimal differentiation has no observable content to theorise about. This is the transcendental domain-specifying claim: observation is differentiation. Established in *Distinction as an Ontological Precondition for Structural Reality* — below this condition, there is no structurally-meaningful reality at all.

Step 2 — Without irreversibility, no records. A reversible distinction can be undrawn. Records in the sense physics requires demand *architectural* rather than merely dynamical stability: their continued availability must be guaranteed by the structure of the admissible dynamics, not left to the luck of which trajectory is realised. A distinction whose re-accessibility is contingent on dynamical accident is not a record; it is a transient pattern that happens to persist. Commitment-level irreversibility is what converts transient patterns into architecturally guaranteed records. This is a stronger condition than the thermodynamic arrow: statistical irreversibility emergent from reversible microdynamics plus initial-condition assumptions (the Past Hypothesis) does not supply the architectural irreversibility records require; it relocates the commitment-level posit from the dynamics' forward evolution to the choice of trajectory, rather than dissolving it.

Step 3 — Without compositional closure, no physics. Records that cannot be composed — that yield no joint fact about the relation between event A and event B — give isolated facts without structure. Physics is a theory of relations among facts; it requires closure under the operations that combine records into further records. Instrumentalist frameworks that deny fact-level closure pay for the denial by losing the ability to state joint facts about distinct events.

Step 4 — Without a minimal architecture, closure fails or trivialises. Compositional closure is not automatic; it requires enough combinatorial and topological structure for the closing operations to yield a stable, non-trivial algebra. Concretely, closure demands three features jointly: non-degenerate combinatorics, non-trivial cycle structure, and stable adjacency. Below a minimum architectural threshold, the combinatorial space is too small to sustain all three simultaneously: closure either fails or trivialises to a single degenerate algebra. The programme's No-Go theorem establishes $K = 7$ as that minimum — the smallest architecture on which non-degenerate combinatorics, non-trivial cycles, and stable adjacency coexist. $K = 7$ is not arbitrary; it is the floor at which closure has room to breathe without collapsing.

Step 5 — Without single-source representation, observable content fragments. A framework whose observable content decomposes into multiple independent functional sources does not have a definite answer to "what are observables functionals of?" — the answer depends on which source is queried. The Single-Source Theorem establishes that any minimal admissible theory expresses its full observable content through a unique functional source: all observables as functionals of the committed record density $\rho(x,t)$, or an equivalent single-source representation.

The result of the admissibility chain. Any admissible theory of the observable must implement differentiation, irreversibility, compositional closure, a minimal architecture (fixed at K = 7), and a single-source functional representation. These five features define a narrow equivalence class of admissible realisations. VERSF is a fully realised instance of that class. The chain does not establish that VERSF is the *unique* realisation in the strict sense; it establishes that admissible realisations share a common structural core, and that core is what the guide calls the fold architecture.

The Elimination Spine (twelve steps)

The admissibility chain is the compressed structural argument. The elimination spine expands it into a twelve-step derivation that walks from the four starting commitments of Part 1 to the fold, ruling out alternatives at each step. The spine does more work than the chain — it does not merely establish that admissible realisations must implement the five features, it establishes specifically that they must contain the 2D intrinsic interface with the (σ, ω) binary data and the $\mathbb{C}^4$ pre-commitment representation.

There are two complementary ways to read the spine:

- A **top-down reading**, starting from the four operational commitments of Part 1.
- A **bottom-up reading**, starting from the Void as a constraint medium (Layer 0).

Both routes must arrive at the same structure. If they do not, the framework fails. The top-down route is formally cleaner — it is how the main Uniqueness paper is organised. The bottom-up route makes inevitability explicit — it shows that each step rules out alternatives rather than merely preferring one. What follows is the combined version: each step can be read either as a consequence of the operational commitments or as a consequence of the Void's constraint structure.

1. Physics requires facts.

Top-down. Without stable, distinguishable, irreversible outcomes accessible to finite observers, there are no experiments, no measurements, no testable predictions. A "theory" producing no facts is not physics.

Bottom-up. A constraint medium that allows reversible distinctions but never produces stable commitments cannot generate observable structure. If no facts can form, no physical theory can be realised.

The two routes agree: facts are the primitive.

2. Facts require finite distinguishability.

Top-down. If arbitrarily fine distinctions could be stably recorded, bounded regions would need infinite information capacity. They don't have it. So there is a minimum resolution below which distinctions cannot be made stable.

Bottom-up. A constraint medium with infinite capacity would never saturate, so commitment would never become costly, so distinctions could always be refined further — and no fact would ever stabilise. Finite distinguishability is not optional; it is what allows commitment to terminate.

Alternatives eliminated: infinite-resolution physics fails on both readings.

3. Facts require irreversibility.

Top-down. If what looks like a fact could be locally undone by a bounded observer, it was a fluctuation, not a fact. Stability, recoverability, and reproducibility all fail simultaneously if local reversibility holds.

Bottom-up. Irreversibility requires that discarded alternatives be physically removed from the local system. If they remain accessible, the system can revert. So commitment requires externalisation of alternatives.

Clarification. The irreversibility required here is *architectural*, not merely dynamical. Architectural irreversibility means the continued availability of a record is guaranteed by the structure of the admissible dynamics — it is structurally unreachable, not just statistically unlikely. Dynamical irreversibility — the familiar thermodynamic arrow — emerges from reversible microdynamics combined with initial-condition assumptions (the H-theorem, the Past Hypothesis). Dynamical irreversibility is real and the programme does not deny it, but it does not supply the architectural irreversibility facts require. A "record" that remains in-principle undoable under the underlying reversible dynamics is a pattern whose continued availability is a matter of dynamical luck rather than architectural guarantee; it is not a fact in the sense the framework uses. Reversible-dynamics-plus-Past-Hypothesis does not refute this step but relocates the commitment-level posit from the dynamics' forward evolution to the choice of which trajectory is actual — which is the posit the step identifies, moved rather than dissolved.

Alternatives eliminated: locally reversible "facts" fail on both readings; merely statistically irreversible "facts" fail the architectural requirement.

4. Externalisation requires a substrate.

Top-down. Once commitment is irreversible, the information about discarded alternatives must go somewhere — the main Uniqueness paper's Theorem 6.2a proves this cannot be absorbed within the accessible sector, so it must be externalised to degrees of freedom inaccessible to bounded observers.

Bottom-up. A system cannot destroy information locally without transferring it to degrees of freedom it no longer controls. Information is not deleted; it is displaced.

Both routes force the existence of the Void. It is not posited; it is the structure irreversibility requires.

5. The Void enforces finite capacity.

Top-down. The inaccessible sector established in step 4 must itself satisfy the operational constraints of $\mathfrak{F}$ — finite observers, bounded resources. A Void with infinite capacity would allow unbounded commitment density in bounded regions, contradicting finite distinguishability.

Bottom-up. If the Void had infinite capacity, commitment would never become costly, and no effective constraints would appear. If it had zero capacity, commitment could not occur at all. The capacity must be finite and nonzero.

Alternatives eliminated: unbounded-capacity and zero-capacity Voids both fail.

6. Capacity limits produce resistance.

Top-down. As commitment density rises toward local capacity, further commitment requires displacing information into an already-loaded Void. The cost of commitment rises.

Bottom-up. Additional commitment becomes progressively harder because discarded alternatives must be absorbed into an already-loaded substrate. Resistance is not imposed — it is the signature of finite capacity encountering load.

This resistance is what underwrites entropy cost, distinguishability limits, and the mechanism of the following step.

7. Resistance defines geometry.

Top-down. Resistance to commitment propagation affects how commitment events distribute and sequence across the substrate. Regions of high load suppress further commitment flow; regions of low load permit it. Crucially, these rate modulations are *directionally structured*: the rate along a given direction depends on the local gradient $\partial_\mu s$ of the coarse-grained entropy field. What finite observers register as distance, duration, and curvature is the structural pattern of these directional modulations. At the level of schematic form (developed in Layer 5), the effective metric reads $g_{\mu\nu}(x) = A(s)\,\eta_{\mu\nu} + B(s)\,\partial_\mu s\,\partial_\nu s$ + higher-derivative corrections, where A and B are functionals of the Void's response (structurally constrained but not yet fully derived — see Layer 5 and Part 6) and $\eta_{\mu\nu}$ is set by the asymmetry between the irreversible time-direction and the reversible space-directions at the fold level. Geometry is not posited but forced *in form*: there is no other way for directional rate modulations to be registered by bounded observers than as an effective pseudo-Riemannian structure of this kind. The explicit functional content of A and B, and the full derivation from effective metric to the Einstein equations, are open.

Bottom-up. Geometry is not a pre-existing manifold on which dynamics happen. It is the Void's response to load gradients, coarse-grained over causal neighbourhoods at scale R^*, and expressed through the schematic metric above in the continuum limit $L \gg R^*$.

Alternatives eliminated: pre-existing geometric substrates are either redundant (they do no work the Void's response does not already do) or inconsistent (they specify structure that the fact distributions cannot determine). A purely discrete relational substrate without the Void's response field fails to produce the coarse-grained directional modulations that an effective metric requires, and therefore fails to recover the weak-gradient reduction to Newtonian gravity in the appropriate regimes.

8. Commitment must occur at a boundary.

Top-down. The transition from reversible to irreversible cannot be global (that would violate local causal boundedness) and cannot be a purely algebraic constraint (a globally invertible projector would not actually produce irreversibility). So it must be a locally defined, geometrically embedded interface separating the two sectors.

Bottom-up. Commitment requires a local separation between reversible possibilities and irreversible facts. This separation must be geometrically realised because the Void's constraint structure is local — there is no global labelling mechanism available.

Alternatives eliminated: non-local commitment mechanisms and bulk-collapse architectures both fail.

9. The interface must be minimal and intrinsic.

Top-down. All candidate interface architectures fail except one:

- Non-local → violates bounded observers (step 8 already eliminated this).
- Extrinsically labelled → commitment depends on description rather than structure, violating admissibility.
- Bulk collapse → redundant under admissibility.
- 1D interface → cannot encode commitment polarity intrinsically (polarity requires reference to an ambient space).
- 3D or higher interface → underdetermined by facts (the same fact distribution is compatible with continuously many inequivalent interfaces).

Bottom-up. The Void's response must localise at a boundary whose structure is uniquely determined by the fact pattern it produces. Only a 2D intrinsic interface has this property.

Only a 2D intrinsic interface survives.

10. The fold is forced.

Top-down. The minimal 2D intrinsic interface carries exactly two independent binary geometric degrees of freedom: commitment polarity $\sigma \in \{0,1\}$ (which side is committed) and reversible orientation $\omega \in \{-1,+1\}$ (direction of traversal). These are independent — neither is a function of the other — and they exhaust the binary structural content of a minimal S^2 surface that can support intrinsic commitment data. Formally: the space of binary geometric functions on the minimal S^2 interface that are admissibility-invariant and support the commitment transition has dimension exactly 2, so it carries $2^2 = 4$ states. No independent third binary function exists on the minimal interface; any candidate third function is either (a) a function of σ and ω (redundant), (b) non-intrinsic (fails admissibility), or (c) requires a larger-than-minimal interface (fails the minimality selection criterion of step 9).

Bottom-up. The Void's local response, reduced to its minimal representational content at an intrinsic 2D commitment boundary, distinguishes exactly the four states $(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$. Any finer response either duplicates information already carried by (σ, ω) or reaches beyond the minimal interface to additional structure that step 9 has already eliminated.

Four states. This is the fold. No smaller structure supports facts; no larger structure is minimal; and the four-state count is not approximate but exact, forced by the dimension of the admissible binary structure on the minimal intrinsic interface.

11. The algebra is fixed.

Top-down. Given a four-state interface with reversible pre-commitment dynamics, the pre-commitment state space is $\mathbb{C}^4$ with unitary evolution. By Hurwitz's theorem the only normed division algebras are $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, $\mathbb{O}$:

- $\mathbb{R}$ fails — the real orthogonal group does not support a global U(1) phase commuting with admissible observables, so relative phases on composite states are not independently accessible and real-linear interference collapses to classical mixture at the level of facts.
- $\mathbb{H}$ is redundant — quaternionic phases produce no new admissible fact outcomes.
- $\mathbb{O}$ is non-associative — cannot support consistent composition of reversible transformations.
- $\mathbb{C}$ satisfies all conditions minimally.

Bottom-up. The Void's reversible response to pre-commitment dynamics must support interference (to distinguish genuine superposition from classical mixture), compositional consistency (associative), and minimality under admissibility. Only $\mathbb{C}$ survives.

Only $\mathbb{C}$ survives.

12. The rest follows.

- $\rho(x,t)$ = load on the Void (the accumulated density of committed facts in each region).
- Entropy = the Void's measure of that load.
- The TPB law = the flow constraint governing how load propagates.
- Geometry = the Void's response to load gradients, producing curvature, gravity, and the effective metric $g_{\mu\nu}$.
- Time = the ordering of commitment events; causality = their directed relations.
- The K=7 substrate architecture, the gauge structure, the specific constants — all downstream consequences of how the fold and its $\mathbb{C}^4$ algebra compose in a Void-constrained substrate.

Convergence. At each step, the alternative is not merely less elegant but structurally inconsistent with finite observers, bounded causality, or irreversible fact formation. The top-down route rules alternatives out on operational grounds; the bottom-up route rules them out on substrate grounds. Both routes converge on the same architecture. The resulting structure is therefore not one possible description of physics among many but sits within a narrow equivalence class of admissible realisations — the class of structures that implement the five features of the admissibility chain and produce the same fact distributions under observational equivalence. VERSF is a fully realised instance of that class. The strict-uniqueness claim (that VERSF is the *only* realisation, not an instance of a pinned class) is stronger than what the current papers establish: *The Layered Foundation of VERSF* and the Uniqueness Master Theorem pin the class, and the Single-Source Theorem pins the functional representation within it, but a proof that the class reduces to a single structure under further admissibility conditions is an open question flagged in Part 6. The correct claim is therefore strong enough to rule alternatives out structurally, and honest enough to say what remains.

A note on the layered structure of the proof. The argument compressed into the spine above is formally a layered proof system in the papers, with each stage delivering an independent result that the next stage builds on. In order: facthood defines the operational content of physics; admissibility (axioms A1–A3 and meta-principles M1–M5) formalises what a fact-producing theory must look like; structural constraints follow from admissibility (finite distinguishability, irreversibility, cyclic substrate, codimension-1 interface, etc.); constraint closure fixes the minimal architecture (the fold with $\mathbb{C}^4$); uniqueness establishes that no alternative minimal architecture exists; saturation (the Saturation Theorem) establishes that the fold exhausts all observable degrees of freedom; and only then do quantitative results (specific constants, the 3/8 commitment barrier coefficient, K=7 geometric consequences) follow from the fixed architecture. The guide compresses these stages for readability, but the papers keep them strictly separated, and each stage can be checked independently.

A note on the relationship between the two routes. The closure result — that any admissible theory producing stable facts must contain the fold architecture — can be derived purely from the operational axioms A1–A3 of the main Uniqueness paper, without invoking the Void as a load-bearing ontological primitive. The closure and saturation papers establish this. The Void, in the guide's presentation, provides the deeper ontological account of *why* those axioms hold — why irreversibility is possible, why capacity is finite, why externalisation must occur somewhere. Readers who prefer a strictly operational derivation can follow the top-down route through the axioms alone; readers who want the physical mechanism follow the bottom-up route through the Void. The framework supports both, and they converge on the same architecture.

This is what makes the framework a narrow-equivalence-class uniqueness claim rather than a proposal among many: an alternative would not just need to differ from VERSF in its surface formulation — it would need to fail either operational physics or substrate consistency, since any admissible alternative is already equivalent to VERSF at the level of fact-supporting structure. Strict uniqueness — reducing the class to a single structure — is a further step on which the programme is in progress; narrow-equivalence-class uniqueness is already established by the current foundational papers.

A note on the two-part structure of the claim. The full uniqueness result has two components that are proved separately. First, *necessity*: any admissible theory producing stable facts must contain the fold. This is the Master Uniqueness Theorem of *Structural Uniqueness of Physical Law from Fact Formation Constraints*. Second, *sufficiency*: once a theory contains the fold, the fold saturates all observable degrees of freedom — there is no residual freedom for alternative physics beyond what the fold generates. This is the Saturation Theorem of *Observable Closure Under the VERSF Fold*. The guide's spine establishes necessity through the 12-step elimination chain; sufficiency is proved by the saturation argument, which operates on a different conceptual axis (observable algebra rather than structural elimination). Necessity says the fold is forced; sufficiency says nothing additional is needed. Together they constitute the complete uniqueness claim.

Whether the argument really closes at every step is the subject of the main Uniqueness paper, its companion repairs paper, and the Saturation paper. Whether VERSF's structure then determines *all* of observable physics — not just the minimal core — is the subject of downstream work on K=7, the Standard Model, and specific constants.

Part 4: Glossary of Key Terms

Several VERSF terms get used in closely related but distinct ways. Pinning them down once saves enormous time later.

Void. The active constraint medium that enforces irreversibility, absorbs discarded alternatives, and bounds the capacity of bounded regions to hold distinguishable structure. Not passive backdrop, not empty space. The Void is what makes commitment events possible (by providing a destination for discarded alternatives), what makes distinguishability finite (by imposing a capacity ceiling), and what produces geometric phenomena including gravity (through its structural response to gradients in commitment density). In the operational language of Part 1, the four starting commitments are the observable face of the Void's constraint structure.

Fold. The 2D commitment interface itself. The specific geometric object — topologically a sphere — that carries the binary degrees of freedom (σ, ω) and separates committed from reversible sectors. "The fold" refers to this interface; "the fold architecture" refers to the full structure of interface plus pre-commitment representation $\mathbb{C}^4$.

Interface. Often used interchangeably with "fold" but more generic. Any admissible commitment mechanism must have an interface; VERSF's specific interface is the fold. When a paper says "interface" it almost always means fold.

Boundary. Usually synonymous with interface or fold. Occasionally refers to the boundary of a specific bounded region (e.g. a causal boundary). Context usually makes it clear.

Commitment event. The irreversible transition from an unresolved alternative to a committed fact. Every observable outcome in physics is a commitment event or a function of commitment events.

Committed record density ρ(x,t). The scalar field describing the spatial and temporal distribution of accumulated committed facts. All observable quantities are ultimately functionals of this field. This is the "single source" in the Single-Source Theorem.

κ-field. A specific scalar field that governs commitment-capacity dynamics. Closely related to ρ(x,t) but not identical — the κ-field is the effective dynamical field, while ρ(x,t) is the density of facts themselves. In some papers these are shown to be structurally equivalent under specific conditions.

Proto-time. The pre-physical ordering of reversible dynamics prior to commitment. Not measurable, not observable, not "time" in the physical sense. Physical time is the ordering of commitment events.

K=7. The constraint dimensionality of stable physical reality — the minimum number of independent binary conditions required for a physical structure to stably and consistently exist. In the substrate realisation, this corresponds to the minimum cell count of the fact-supporting simplicial architecture (six hexagonal-boundary sites plus one hub). Pinned by two convergent structural results: the No-Go theorem on the minimal fact-supporting substrate, and the Convergence Theorem on six labelled constraints reducing to four mathematical structures, with the Diophantine equation $2^K + 8 = N(N-1)/2$ (unique solution $K = 7$, $N = 17$) providing a fully unconditional integer-exact anchor. The symmetry group is $GL(3,2) \cong PSL(2,7)$, the automorphism group of the Fano plane. Layer 6 develops the full structure.

BCB — Bit Conservation and Balance. The principle that information (bits) is conserved across commitment events, with balanced accounting between committed and discarded alternatives. One of the core regulative principles of the framework.

TPB — Ticks-Per-Bit. The principle that the rate of commitment events is bounded, relating the temporal granularity of fact production to information-theoretic constraints.

VERSF. The umbrella framework — Void Energy–Regulated Space Framework — that includes BCB, TPB, and the fold architecture as interconnected components. "VERSF" usually refers to the whole programme; "the fold" refers to the specific minimal structure within VERSF.

PAR / CC. Pre-Factual Algebraic Reversibility and Compositional Completeness. The formal properties that the pre-commitment sector must satisfy. These are derived conditions, not assumptions.

IAC — Internal Admissible Closure. The structural requirement that the reversible pre-commitment sector must be closed under its own admissible operations. Related to the PAR/CC conditions.

DCPC — Delayed-Commit Pattern Computing. A computational architecture inspired by VERSF, in which commitment is deferred as long as possible. Applied primarily in the quantum computing papers.

ξ — Closure scale / mesoscopic coherence scale. A specific length scale, approximately 8.2×10^{-5} metres, derived from VERSF and appearing in multiple phenomenological contexts (cosmology, condensed matter, biology). It marks the transition between microscopic commitment dynamics and macroscopic classical behaviour.

Distinguishability. Used technically: the minimum resolution δ_{min} below which outcomes cannot be stably separated. Finite distinguishability is one of the core starting commitments.

Admissibility. A structure or description is "admissible" if it satisfies the minimal conditions for doing operational physics (finite observers, bounded resources, stable facts). The admissible theory class $\mathfrak{F}$ is the class of theories in which operational physics is possible.

Minimality. A structure S is minimal within the admissible class if no proper substructure of S suffices to produce the admissible facts that S produces. Equivalently: if you remove any piece of S and the remainder still accounts for all admissible facts, S was not minimal. Minimality is the selection principle under which the guide rules out bulk-collapse architectures, quaternionic extensions, higher-genus interfaces, and similar "admissible but non-minimal" alternatives. It is not an aesthetic preference; it is the structural content of condition 6 (admissibility of description): any surplus structure that produces no new facts is eliminable by observational equivalence, so the minimal admissible structure is the canonical form.

Observable closure. The property that the fold's structure already determines all admissible observables, leaving no room for additional independent observable degrees of freedom. This is the content of the Saturation Theorem.

Part 5: Reading Paths

The programme is too large to read linearly. Here are suggested sequences for different entry points.

A bibliographic note. Several papers cited in this guide are internal drafts not yet publicly available at versf-eos.com. They are marked below with "(in preparation)." As of the current version of the guide, these are: *Structural Uniqueness of Physical Law from Fact Formation Constraints*, *Structural Repairs to the VERSF Uniqueness Programme*, *Observable Closure Under the VERSF Fold: A Saturation Theorem*, *A No-Go Theorem for Non-Simplicial Relational Substrates — K=7* (the standalone K=7 no-go paper, distinct from *K=7: The Constraint Dimensionality of Stable Physical Reality* which is on the site), *Two Kinds of Time: Proto-Time and Physical Time*, and *Single-Source Theorem*. Readers who want to access these papers should contact the author directly. All other cited papers are retrievable from versf-eos.com. Elsewhere in the guide, these papers are cited without the "(in preparation)" marker for readability; this note is the canonical list.

Path A: The Structural Uniqueness Argument (for mathematical physicists)

If you want the core claim — that the fold is the unique minimal architecture for any fact-producing universe, up to observational equivalence — read these in order:

1. *Structural Uniqueness of Physical Law from Fact Formation Constraints* (the main Uniqueness paper) — **in preparation**
2. *Structural Repairs to the VERSF Uniqueness Programme* (the companion repairs paper) — **in preparation**
3. *Observable Closure Under the VERSF Fold: A Saturation Theorem* (the sufficiency paper) — **in preparation**
4. *A No-Go Theorem for Non-Simplicial Relational Substrates — K=7* — **in preparation**
5. *Complex Hilbert Space from Distinguishability*

Stop there. These five papers carry the core structural argument. Everything else is either foundational background or downstream phenomenology.

Path B: The Conceptual Framework (for philosophers of physics)

If you care about what VERSF *means* rather than the technical proofs:

1. *Why Physics Needs Facts to Exist at All*
2. *The Hidden Conditions That Make Physics Possible*
3. *Two Kinds of Time: Proto-Time and Physical Time* — **in preparation**
4. *The Void Is Not Nothing*
5. *Measurement as Commitment: Why Quantum Systems Are Relational*
6. *Where Time Really Comes From*
7. *The Geometry of Facts: Why Reality Needs Two Dimensions First*

Path C: Quantum Foundations (for quantum mechanics specialists)

1. *Complex Hilbert Space from Distinguishability*
2. *Quantum Mechanics as Admissibility Fixed Point*
3. *Why a Fact-Producing Universe Must Satisfy Interference, Isotropy, and Representational Invariance*
4. *Why Quantum Systems Multiply Instead of Add: A Physical Origin of the Tensor Product*
5. *Measurement as Commitment*
6. *Wave-Particle Duality as Pre-Temporal Commitment Dynamics*
7. *The Quantum Zeno Effect as Suppression of Irreversible Commitment*

Path D: Cosmology and Gravity (for cosmologists)

1. *Record-Theoretic Emergence of Spacetime Geometry and Gravitational Dynamics*
2. *Gravity from Fold Density Gradients*
3. *Why Emergent Gravity Must Be Spin-2*
4. *Deriving the Cosmological Constant*
5. *Tensor Perturbations in the VERSF Framework* (includes $r \approx 0.027$–0.033 prediction for LiteBIRD)

6. *Where the 1/4 in Black Hole Entropy Comes From*

Path E: Numerical Predictions and Constants (for phenomenologists)

1. *From Admissibility to 1/137: The Capstone of the VERSF Programme*
2. *From Structure to Constant: A Conditional Derivation of the Fine-Structure Constant*
3. *Deriving the Cosmological Constant*
4. *The Standard Model from Hexagonal Geometry*
5. *Tensor Perturbations in the VERSF Framework*
6. *Why the Universe Has a Smallest Length — and Why Gravity Is So Weak* (Planck-scale derivation)

Path F: General Reader

Start with any paper whose title begins with "Why" or "How" — these are written for non-specialists. Suggested order:

1. *Why Physics Needs Facts to Exist at All*
2. *Where Time Really Comes From*
3. *The Void Is Not Nothing*
4. *How Physical Facts Come Into Existence*
5. *Why Something Rather Than Nothing? Dissolving Physics' Deepest Question*
6. *Rethinking the Big Bang: A Universe Built from Information, Not Spacetime*

How VERSF Differs from Cognate Programmes

A sophisticated reader encountering VERSF may notice structural similarities to several existing programmes in foundational physics. The similarities are real; the differentiations are important. *The Layered Foundation of VERSF* (§6) positions VERSF explicitly against five cognate programmes, and the summary here follows that treatment.

Quantum reconstruction programmes (Hardy, Chiribella et al., Masanes–Müller)

Quantum reconstructions typically start from axioms about probabilities, information capacity, or operational distinguishability, and derive the Hilbert-space formalism of quantum mechanics as the unique structure satisfying those axioms. They answer the question: given that we observe probabilities behaving a certain way, what mathematical structure must underlie those probabilities?

VERSF starts one level deeper. It does not take probabilities as its starting point. It takes *facts* — stable, recoverable, reproducible, irreversible outcomes — as the primitive. Probability structure, interference, Hilbert spaces, and the Born rule all emerge as consequences of the requirements that facts must satisfy. Three concrete differences:

Starting point. Quantum reconstructions start from probability distributions over outcomes. VERSF starts from the irreversibility of outcome formation. Irreversibility is not derivable from probability axioms alone; in VERSF it is a primitive requirement.

Geometric content. Quantum reconstructions derive Hilbert space but are generally silent on the geometric structure from which it emerges. VERSF derives a specific two-dimensional commitment interface with binary geometric data (σ, ω) as the physical object the algebraic structure describes.

Downstream commitments. Quantum reconstructions typically stop at quantum mechanics. VERSF continues through K=7 substrate architecture, emergent spacetime geometry, specific numerical predictions (α^{-1}, Λ, r), and the gauge structure of the Standard Model.

A concrete example: Hardy's 2001 axiomatisation postulates continuity as Axiom 5 — that any two pure states can be connected by a continuous reversible transformation. In Hardy's system this is a fundamental assumption with no deeper justification. In VERSF the analogous statement emerges as a consequence of reversible pre-commitment dynamics on the fold: the pre-commitment sector evolves unitarily by construction, so continuous reversible connectivity of pure states is a structural output rather than an input. This is the typical pattern — what quantum reconstructions assume operationally, VERSF derives from the fold architecture plus the Void's constraint structure.

Hardy is a methodological cousin of VERSF, not a foundational one: the alliance is on how to derive, not on what the primitives are.

Wheeler's "it from bit"

Wheeler's programme is closest to VERSF's operational spirit. Both take information-theoretic preconditions on observation as generative. The difference is ontological: Wheeler's bit is *ontologically strong* — the bit is posited as the fundamental. VERSF does not treat distinguishability as the ontic fundamental; it treats it as the *epistemic primitive*, with potentiality (the Void) as its ontic complement. The programmes agree that information-theoretic preconditions are generative, and disagree on whether those preconditions are themselves the bedrock of being. VERSF's two-primitive ontic level is the load-bearing distinction.

Rovelli's relational quantum mechanics

Rovelli's programme also treats actualisation as event-like, and Rovelli's "observers" are any physical systems rather than agents with minds — so the superficial similarity to VERSF's commitment events is real. The substantive disagreements are two:

First, whether actualisation is relational. For VERSF a commitment is fact *simpliciter* — a fact in its own right. For Rovelli every fact is fact-relative-to-a-system, and no fact can stand independently of the system to which it is relativised. This is not a matter of framing; it is a substantive disagreement about whether commitments have standalone ontological status.

Second, whether there is a substrate aspect of existence. VERSF affirms the Void as non-relational substrate. Rovelli denies any non-relational ontology. The two programmes share the conviction that observation-like events are foundational; they differ on non-relational facts and on substrate.

Zurek's existential interpretation and einselection

Zurek derives distinguishability from quantum dynamics: environment-induced superselection picks out preferred states as an emergent consequence of decoherence. VERSF runs in the opposite direction — distinguishability is the precondition, not the output; commitment dynamics are derived as consequences of the primitive rather than emerging from an underlying wavefunction. The two programmes are duals on the explanatory arrow.

Scientific realists (Maudlin, Norsen, Ladyman–Ross)

Realists press the question of primitive ontology: what are commitment events — points, fields, particles, events at locations? VERSF's answer distinguishes it from both primitive-ontology realism (Maudlin) and eliminativist structural realism (Ladyman–Ross). Commitment events are *not* located at pre-given spatial points, because space is itself derived from the pattern of commitments; but they are *not* merely structural either, because their content — spectral density, barrier height, architectural role, combinatorial architecture — is theoretically specified rather than operationally defined. The framework is closer to a moderate structural realism: structures are real, but the relata are not eliminated — they are the commitments themselves, with specified content.

The programme's distinctive position

VERSF's distinctive position across these comparisons is defined by four features taken together: (i) a two-primitive ontic level with an explicit substrate aspect, (ii) non-relational actualisation — commitment as fact simpliciter rather than system-relative fact, (iii) commitment treated as a physical source entering the dynamics, and (iv) an epistemic primitive with substantive rather than merely formal structural content. Each cognate programme occupies pieces of this region; none occupies the whole of it.

This is not a claim that VERSF is superior to any of the cognate programmes. They answer different questions, and a complete picture of foundational physics may eventually require more than one of them. But confusing VERSF with any of them — in particular, dismissing it as "another quantum reconstruction with extra stuff," "it from bit restated," or "Rovelli with a Void bolted on" — misses what makes it distinctive.

Part 6: What Is Still Open

No research programme is complete, and VERSF is honest about its remaining work. The sections below give an accurate accounting of where things stand, organised into three tiers — established, partially derived, and open — followed by a detailed list of specific technical exposures.

Established

Results at this level are either derived from the axioms without open steps, or proved via integer-exact unconditional arguments, or established with explicit calculations whose correctness is independently verifiable.

Structural core. The structural uniqueness of the fold architecture within the admissible theory class $\mathfrak{F}$ is established by the main Uniqueness paper, subject to the five repairs in the companion paper. Given the four starting commitments, any theory capable of producing stable facts for finite observers must contain a structure equivalent to the fold: a 2D intrinsic commitment interface with binary data (σ, ω) and a reversible pre-commitment representation over $\mathbb{C}^4$. The observable closure result — that once a theory contains the fold, all admissible observables are already determined — is established in the Saturation paper, via characterisation of the canonical observable algebra as $M_4(\mathbb{C})$ and explicit killing of the three remaining escape routes. The Single-Source Theorem, that all observables are functionals of ρ(x,t), is derived via the Saturation argument combined with admissibility equivalence.

Substrate architecture. K = 7 as the constraint dimensionality of stable physical reality is established by convergence of four distinct mathematical structures on a single integer, with C5 — the Diophantine equation $2^K + 8 = N(N-1)/2$ with unique solution (K, N) = (7, 17) — providing a fully unconditional integer-exact anchor. The K = 7 No-Go theorem on the minimal fact-supporting substrate is established within its stated structural conditions. The reduction of "six labelled constraints" to "four mathematical structures" under honest independence analysis is explicit. The rank-6 / nullity-1 structure of the hexagonal cell's constraint matrix (C1 of the K=7 paper) is established via the graph-Laplacian argument on the hub-and-boundary graph G_7 in *Completing the Interface Bridge*, together with the Honeycomb Theorem and the orbit-stabiliser enumeration showing that the hexagonal cell admits exactly 7 independent constraint support sites (6 boundary vertices + 1 hub) — exclusion of K > 7 is therefore proved, not argued. The six-channel interface structure, the uniform 1/6 allocation across channels, the 1/6 inverse participation ratio, and the exclusion of cross-channel covariance terms from the primitive local observable algebra are also established in the same paper from five primitive postulates plus the Economy Axiom.

Causal structure. The existence of a finite invariant propagation speed c is derived from finite distinguishability, bounded local update, and causal consistency — not postulated. $c_T = 1$ (gravitational waves travel at the speed of light, consistent with GW170817) is proved. The emergence of time, causality, and effective spacetime from the commitment structure is established in the relevant papers.

Geometry (structural content). The transition from discrete commitment events to continuous spacetime geometry is established via coarse-graining at the fundamental scale R^* and the Void's response to entropy gradients. The Poisson equation for gravitational potential is derived in the weak-gradient limit as the unique admissible form in the class of local, isotropic, linear, second-order, scale-free field equations. The inverse-square law is derived as the unique admissible vacuum solution in the same class.

Ontological clarifications. Renormalisation depth is established as a descriptive parameter rather than a spatial direction, via a formal no-go theorem on non-injective coarse-graining semigroups. The distinction between descriptive and emergent geometry is formal, not merely interpretive. The bulk geometric interpretation in holographic constructions is established as emergent and regime-dependent, not primitive.

Numerical anchors. The bare fine-structure coupling $\alpha^{-1}\_\text{bare} = 2^7 \cdot 15/14 \approx 137.143$ is derived with no parameter fitting. The proton shell count N = 17 is derived from BCB closure constraints (not fitted) and matches the N appearing in the C5 Diophantine equation. The κ-field mass coefficient √(4/3) from PGL(3,2) irreducibility on V_6 is derived.

Partially Derived

Results at this level have the K = 7 structure entering correctly, with the leading-order or structural term established, and with specific calculational or normalisation steps still required for completeness.

The fine-structure constant correction. The structural origin of the IPR correction to α^{-1}_bare is identified, and the fact that it arises from hexagonal-isotropy symmetry on six boundary channels is derived. The specific correction value h ≈ 0.010 has not yet been derived from first principles; the current best candidate h = 1/C(12,2) = 1/66 gives $\alpha^{-1} \approx 136.994$, bracketing the target 137.036 from below.

The cosmological constant. The functional form ξ(K) ∝ exp(π · 2^K / 2) together with $\Lambda = 3/\xi^2$ places K = 7 uniquely in the observed cosmological band, with adjacent-integer separation of ~53 orders of magnitude making the selection robust. The absolute proportionality constant requires the normalisation derivation in *Geometric Closure*, which closes the ~52-order gap between nominal and observed Λ.

Black hole entropy. Area-linear scaling S ∝ A is derived, and K = 7 enters the coefficient consistently. The exact 1/4 numerical coefficient requires the normalisation derivation that converts between the VERSF discrete constraint count and the continuous entropy measure in Planck units (developed in *Completing the Interface Bridge*). The VERSF count currently gives a coefficient too large by ~19; the normalisation closes this gap.

The metric schematic. The form g_μν(x) = A(s) η_μν + B(s) ∂_μ s ∂_ν s + ... is structurally correct, with Lorentzian signature inherited from the fold's time/space asymmetry. The explicit functional forms of A(s) and B(s) from K = 7 closure and TPB dynamics are asserted structurally rather than derived; explicit derivation is an open task.

Tensor gravity. The structural argument for the promotion of the scalar potential to a rank-2 symmetric field is in place. The detailed derivation of the linearised Einstein equations from the commitment dynamics, and the reduction to full nonlinear Einstein equations $G_{\mu\nu} = f(\rho, \partial\rho, \partial^2\rho, ...)$, are in progress.

The tensor-to-scalar ratio. $r \approx 0.027–0.033$ is derived from the κ-field spectrum; the range collapses to a single value once the Fano-plane GL(3,2) correction is computed and the scalar stiffness Q_S is determined. LiteBIRD (~2032) will test this range.

The coherence scale. The functional form of $\xi(K)$ is derived; the proportionality constant that sets the absolute value is the subject of the *Geometric Closure* companion paper.