The Hierarchy Problem as a Category Error

A VERSF Interpretation of Scale Separation, Emergent Mass, and Closure Dynamics

Keith Taylor *VERSF Theoretical Physics Programme*

Plain-Language Summary

For nearly fifty years, physicists have been puzzled by the size of the Higgs boson — the particle that gives other particles their mass. The Higgs is far, far lighter than the simplest calculations say it should be. The gap is staggering: about seventeen zeros' worth of difference between the prediction and the measurement. Under the standard understanding, keeping the Higgs this light would require nature to perform an absurdly precise cancellation, an arithmetic of opposing numbers that has to come out balanced to better than one part in a hundred million trillion trillion. That much precision feels suspicious — like a coin landing on its edge a thousand times in a row — and so for half a century most theoretical physics has been spent looking for hidden ingredients (extra particles, extra dimensions, new symmetries) that could explain why the cancellation happens.

This paper argues that the puzzle has been misdiagnosed.

The trouble starts with what is being compared. The Higgs's mass is one number; the "Planck scale" — the scale at which gravity itself becomes a quantum thing — is another. The standard story treats these as the same kind of number, separated by a vast distance that ought to be explained. This paper says they are not the same kind of number at all. Comparing them is like asking why middle C on a violin is so much lower in pitch than the chemical bonds holding the violin's wood together.

That comparison sounds odd — and that's the point. A note like middle C is a slow, gentle vibration — a few hundred wiggles per second. The chemical bonds in the violin's wood vibrate vastly faster, trillions of times a second. The gap between them is enormous. But nobody finds it puzzling. Nobody demands a precise cancellation to explain why middle C is "unnaturally low" compared to bond energies. We just understand that a musical note is not the same kind of thing as a chemical bond. The note is what happens when the whole violin vibrates as one. The bonds are what hold the wood together so that vibration is possible at all. Both are real, both matter, and they are deeply connected — without the bonds, no violin, no note. But they live at different levels of reality, and there is no mystery in their being wildly different in size.

This paper argues that the Higgs is to the deep structure of the universe what middle C is to the violin. It is what emerges when a vast amount of underlying activity cooperates as one. The Planck scale, on the other hand, is not a place where heavier particles live, waiting to be discovered. It is something more fundamental — a boundary on what reality is able to distinguish at all. Within the framework this paper uses (called VERSF, which builds space, time, and quantum behaviour out of more basic ideas about how things can be different from each other), these two scales belong to different layers of reality. The deep substrate shapes the Higgs the way wood shapes a violin's tone, but it does not make the Higgs heavy any more than the wood makes middle C high-pitched. The huge gap between the two scales is no longer a crisis. It is what one expects from a reality built in layers.

There is also a quantitative payoff. When the framework's own building blocks are combined in a natural way, they predict that any genuinely new physics should appear, if at all, somewhere around a few thousand times the energy of the Higgs itself — broadly the range that the Large Hadron Collider in Geneva has been searching for over a decade. Other proposed solutions to the hierarchy puzzle expected a whole crowd of new particles in this range. The Collider has found none. This framework expects exactly that: nothing dramatic, just a gradual softening as one approaches the band edge. The same architecture also reproduces, to within a percent, an unrelated measurement from the cosmic microwave background — the faint glow left over from the early universe. Two pieces of evidence from one underlying picture.

The paper does not claim that mainstream physics is wrong in general. It claims only that one specific habit of thought — extending equations developed for our familiar level of reality all the way down to the deepest underlying structure — crosses a boundary it was never built to cross. The puzzle of the Higgs's lightness is not, on this reading, a failure of nature to protect it. It is evidence that the deepest scale of reality is simply not part of the same world the Higgs lives in.

Readers who prefer the formal statement can proceed directly to the technical abstract below. Readers who want the conceptual core without equations may skip to §7 (*Why This Is Not Composite Higgs*), §10 (*Why Naturalness Looked Compelling*), and §14 (*Conclusion*), which together convey the principal claims of the paper while remaining largely equation-free.

Table of Contents

Abstract

The hierarchy problem of contemporary particle physics is conventionally framed as a naturalness crisis: the electroweak (EW) scale lies roughly seventeen orders of magnitude below the Planck scale, and a fundamental scalar Higgs field should — barring delicate cancellations or new ultraviolet (UV) symmetries — receive quadratic radiative corrections that drive its mass toward the cutoff. Decades of beyond-Standard-Model construction have proceeded under this diagnosis.

This paper argues that the diagnosis itself rests on a category error. Within the Void Energy–Regulated Space Framework (VERSF), the Planck scale and the Higgs scale do not occupy the same ontological stratum. The Planck scale is a **closure threshold** of the substrate — a constitutive limit on admissible distinguishability — whereas the electroweak scale is an **emergent coherence scale** of the record field built on that substrate. We formalise the distinction as the *Ontological Separation Principle*, and the impossibility of treating closure-layer structure as propagating field excitations as the *Closure Non-Propagation Principle*. Together these dissolve the premise that the renormalisation group can extrapolate scalar mass corrections from the EW scale to the Planck scale at all.

We then construct a concrete admissibility suppression kernel F_adm(k) — the Fourier-side counterpart of the VERSF-derived memory kernel K(τ) — that regulates the scalar two-point

function without invoking new symmetries. The effective record-layer closure scale k_C is identified with the κ-field mass m_κ, which is itself a theorem of the K=7 minimal fact architecture and the Causal–Coherence Compatibility (CCC) threshold: $m^2_\kappa = \lambda_{eff} \cdot \xi^{-2} = (4/3) \cdot \xi^{-2}$, where $\xi = (\hbar c/\rho)^{1/4}$ is the CCC coherence scale and $\lambda_{eff} = 4/3$ is the lowest eigenvalue of the closure operator on the physical state space. The hierarchy $k_C / M_P = \sqrt{(4/3)} \cdot \ell_P / \xi$ is therefore the *form* of a derived dimensionless ratio rather than a fine-tuning cancellation. The numerical value depends on which ξ governs the electroweak sector — an identification that has not yet been derived from the master action; §4.3 sets out three structural responses and is explicit about which steps are theorems and which remain conditional. Among these, one response makes direct numerical contact with observation: the geometric-mean rung $E_{geo} = \sqrt{(M_P \cdot m_\kappa)} \approx 5$ TeV, built from M_P and the observed cosmological vacuum-energy density (via the CCC threshold and the K=7 spectral coefficient $\sqrt{(4/3)}$), places v = 246 GeV within one order of magnitude of the predicted band. This is a one-parameter relation between three observed scales (ρ_Λ, M_P, v) through a structural coefficient derived from the substrate architecture — not a first-principles derivation of v, but not a numerological coincidence either, since the geometric-mean form has been noted previously in the cosmological-constant-as-IR-cutoff literature (Banks; Cohen, Kaplan, Nelson) and the VERSF reading supplies its structural origin. The smooth-turn-off phenomenology that follows is qualitatively consistent with the LHC's null results on TeV-scale resonances. We argue that the physical interpretation of the standard quadratic divergence is regulator-dependent and not scheme-independent; supply a VERSF-specific coherence functional containing record-current transport and an admissibility term; reframe the UV–IR geometric-mean relation as evidence for a structural coherence *band* rather than a numerical coincidence; and explain why the conventional naturalness argument looked compelling for so long without being correct in the case at hand. The composite-Higgs reading is distinguished explicitly from the VERSF reading. The closing claim is direct: the hierarchy problem is not a failure of nature to protect the Higgs — it is evidence that the Planck scale is not part of the Higgs sector at all.

---

1. The Conventional Statement and Its Hidden Premise

The hierarchy problem can be stated tightly. The electroweak vacuum expectation value $v \approx 246$ GeV sets the Higgs mass scale $m_H \approx 125$ GeV. The reduced Planck mass is $M_P \approx 2.4 \times 10^{18}$ GeV. Their ratio is

$M_P / v \approx 10^{16}$,

and the squared ratio — which controls quadratic radiative sensitivity — is roughly 10^{32}.

In Wilsonian effective field theory, a fundamental scalar coupled to states up to a cutoff Λ receives a one-loop mass correction of the schematic form

$\delta m_H^2 \sim (c / 16\pi^2) \Lambda^2$, (1) with c an order-unity combination of couplings. If Λ is identified with M_P, the bare and counterterm contributions to m_H^2 must cancel to about one part in 10^{32} — the familiar fine-tuning indictment.

The diagnosis carries three premises, usually unstated:

(P1) The cutoff Λ refers to a *physical* sector of degrees of freedom that the Higgs couples to. **(P2)** The Planck scale *is* such a sector (a UV completion populated by particles, strings, branes, or other field-theoretic excitations). **(P3)** The Higgs field is *fundamental* — an irreducible scalar degree of freedom of the same ontological kind as the substrate at M_P.

The standard programme (supersymmetry, composite Higgs, extra dimensions, conformal extensions) accepts (P1)–(P3) and seeks a symmetry that protects m_H against the implied corrections. VERSF rejects (P2), re-examines (P3), and further argues that the physical interpretation of (1) is regulator-dependent rather than scheme-independent. The result is not a new symmetry but a reclassification of what the two scales *are*, and a derivation of why the extrapolation that produces the apparent crisis is physically meaningless beyond the closure threshold even where it remains a well-defined formal expression.

Two empirical contacts from one architecture. The argument developed here is structural, but the same architecture has two independent empirical contacts that a reader may want to know up front. The K=7 minimal-fact spectral data that fixes the $\sqrt{(4/3)}$ coefficient in the κ-field mass also yields the CMB scalar spectral index $n_s \approx 1 - 2/N\bigstar \approx 0.964$ at $N\bigstar = 55$ e-folds, compared with the Planck 2018 measurement $n_s = 0.965 \pm 0.004$ (Taylor, *The Spectral Density of the Commitment-Event Bath in the VERSF Framework*) — a percent-level agreement on an independently measured cosmological observable. Combining the same κ-field mass with the Planck scale via the closure–CCC geometric mean (§8) gives an emergent coherence-band edge at $E_{geo} \approx 5$–6 TeV, placing $v = 246$ GeV within one order of magnitude of the predicted band and consistent with the LHC's null results on TeV-scale new physics. Two empirical contacts — a percent-level CMB observable and a TeV-scale coherence-band edge — issue from one substrate architecture. The structural argument below stands on its own, but a reader inclined to grant the rest of the paper more rope after seeing the empirical anchors has them here.

---

2. The Two-Layer Ontology of VERSF

VERSF is built on two primitives: **finite distinguishability** and **irreversible commitment**. Physical reality emerges from sequential commitment events constrained by the Physical Admissibility Framework (PAF), Bit Conservation and Balance (BCB), and Ticks-Per-Bit (TPB) bookkeeping. From these primitives, two structurally distinct layers arise.

Closure layer $\mathcal{L}_C$. The substrate of admissible distinguishability. The Planck length $\ell_P = \sqrt{(\hbar G/c^3)}$ and Planck mass M_P are not characteristic energies of any propagating excitations on this layer. They are *closure thresholds*: bounds on the minimum admissible resolution of commitment events. The Two-Planck framework formalises this — ℓ_P bounds spatial admissibility, while a complementary temporal closure scale bounds the TPB rate.

The closure layer should not be interpreted as a hidden spacetime populated by smaller objects — there is no microscopic ether of Planckian particles waiting to be uncovered. Spacetime itself is emergent only within the record layer; $\mathcal{L}_C$ is the set of admissibility conditions on commitment events from which spacetime structure subsequently emerges. Universal gravitational coupling is preserved by this reading because gravity arises statistically from commitment-density distributions in the record layer (Gravity from Tensorial Closure of Record Dynamics) rather than from propagating Planckian quanta. The closure layer conditions gravitational dynamics without entering them as a population of mediator states.

Record layer $\mathcal{L}_R$. The emergent field-theoretic structure built from coherent commitment patterns above the closure threshold. Standard Model fields, including the Higgs, live here as coherence eigenmodes of the record field. The record layer admits a Lagrangian description, perturbative expansions, and renormalisation-group flow.

The layered structure is summarised in Figure 1.

Figure 1. *Layered ontology and the locus of phenomena.*

```
┌──────────────────────────────────────────────────┐
│ CLOSURE LAYER  𝓛_C                               │
│ ─────────────────────                            │
│ • Finite distinguishability                      │
│ • TPB / BCB bookkeeping                          │
│ • PAF admissibility                              │
│ • Planck thresholds ℓ_P, M_P                     │
│ • NO propagation, NO field excitations           │
└──────────────────────────────────────────────────┘
                    │  emergence transition
                    │  (admissibility closure  𝒜[ρ]=0)
                    ▼
┌──────────────────────────────────────────────────┐
│ RECORD LAYER  𝓛_R                                │
│ ─────────────────────                            │
│ • Commitment density ρ, current J_ρ              │
│ • Gauge & matter fields                          │
│ • Coherence functional 𝒞[ρ]                      │
│ • Higgs as coherence eigenmode of ρ              │
│ • RG flow valid WITHIN this layer only           │
└──────────────────────────────────────────────────┘
                    │  low-energy projection
                    ▼
┌──────────────────────────────────────────────────┐
│ STANDARD MODEL OBSERVABLES                       │
│ ─────────────────────                            │
│ • v ≈ 246 GeV,  m_H ≈ 125 GeV                    │
│ • Gauge group SU(3)×SU(2)×U(1)                   │
│ • Fermion masses, CKM/PMNS mixing                │
└──────────────────────────────────────────────────┘
```

The category error of the standard hierarchy problem is now visible: premise (P1) assumes Λ couples to record-layer degrees of freedom up to and *through* the Planck scale. But the Planck scale is not a record-layer scale at all — it is a closure threshold belonging to $\mathcal{L}_C$ beneath $\mathcal{L}_R$.

2.1 The Ontological Separation Principle

We state the OSP first, although its derivation depends on the CNPP (§3) and the admissibility kernel (§4), which together supply the inputs from which $\mathscr{R}_{RG}$ is constructed; the linear order chosen here is expository rather than logical.[^principle-note]

[^principle-note]: "Principle" is used throughout this paper in the sense of a named structural claim derived from VERSF primitives, not in the sense of an unargued axiom. Both the OSP below and the CNPP of §3 are derived from inputs identified in §4 and §3 respectively.

The renormalisation-group operator $\mathscr{R}_{RG}$ of an effective field theory is not a free abstract construction. It is built from the β-functions of the theory, which in turn are computed from loop integrals over the propagating degrees of freedom of the record layer. Under VERSF, those loop integrals are weighted by the admissibility kernel $F_{adm}(k)$ introduced in §4: the inputs to β_m are not arbitrary momenta but admissibility-restricted momenta. We can therefore state OSP as a *consequence* of how $\mathscr{R}_{RG}$ is constructed:

Ontological Separation Principle (OSP). The record-layer RG operator $\mathscr{R}_{RG}$ is constructed from loop integrals weighted by the admissibility kernel $F_{adm}(k)$, where

$$F_{adm}(k) \to 0 \text{ as } k \to k_C. \quad (2)$$

Consequently the β-functions that generate $\mathscr{R}_{RG}$ receive vanishing contributions from momenta above k_C, and $\mathscr{R}_{RG}$ is well-defined only on configurations with characteristic momenta below k_C:

$$\mathscr{R}_{RG} : \mathcal{L}_R(\mu < k_C) \longrightarrow \mathcal{L}_R(\mu < k_C). \quad (3)$$

The integration

$$m_H^2(M_P) = m_H^2(v) + \int_v^{M_P} \beta_m(\mu)\, d \ln \mu \quad (4)$$

exists as a formal Riemann integral, but the integrand $\beta_m(\mu)$ is itself defined through inputs that vanish for $\mu > k_C$. The integral therefore receives no physical contribution from $\mu \in (k_C, M_P)$, and its physical content is captured by the intra-record-layer flow up to k_C. Reading (4) as carrying Planck-scale sensitivity is inconsistent with the admissibility regulation that defines its integrand in the first place.

The OSP is not a denial of renormalisation-group reasoning. It is a statement that the domain of validity of $\mathscr{R}_{RG}$ is bounded above by the closure threshold of $\mathcal{L}_R$, and that crossing that threshold is not a change of scale within a single theory but a change of theory itself. Equation (4), read naively to M_P, mistakes a categorical boundary for a regular point in parameter space

— the integrand exists as a mathematical object but no longer represents the physics it was constructed to represent.

The value of k_C for the EW sector. The OSP statement is option-neutral with respect to which substrate scale plays the role of k_C. Under the leading reading of §4.3 (option (iii)), the operational k_C for the EW sector is the geometric-mean rung $E_{geo} \approx 5$–6 TeV, *not* the cosmological $m_\kappa \approx 2.8$ meV; under the alternative resolutions (i)/(ii) k_C is m_κ at a sector-specific ξ to be derived. In all cases the structural form of (2)–(4) is preserved; what differs across the readings is the numerical scale at which RG flow within $\mathcal{L}_R$ terminates. A reader following the leading option (iii) thread should read the k_C of this subsection as the few-TeV scale, consistent with successful EFT practice up to LHC energies.

---

3. The Closure Non-Propagation Principle

The OSP is grounded in the admissibility kernel, but the kernel's existence raises a prior question: *why* should closure-layer structure not propagate within the record layer to begin with? The answer follows from the VERSF primitives.

A propagating excitation in $\mathcal{L}_R$ requires, minimally, three conditions. These are stated as VERSF-native characterisations, but each aligns with an independently motivated requirement of standard quantum field theory — supplying cross-validation rather than tautology:

(C1) Sequential state distinguishability. Successive states along the excitation's worldline must be distinguishable; otherwise no propagation has occurred. In QFT terms: the theory must admit a pole structure with distinct in- and out-states for the propagating mode.

(C2) Admissible TPB transport. The Ticks-Per-Bit bookkeeping must support sequential commitment events at the propagation rate. In QFT terms: unitarity. Probability is conserved because the commitment budget is conserved.

(C3) Stable record continuity. Commitment events must form a coherent record sustaining the excitation's identity across ticks. In QFT terms: locality of correlation functions. Records cluster in time and space because commitment events do.

The fact that VERSF-native characterisations (C1)–(C3) coincide with the standard QFT requirements of pole structure, unitarity, and locality is itself substantive: it would not be coincidence if VERSF correctly identifies the substrate from which field-theoretic dynamics emerge.

Status of the (C1)–(C3) ↔ QFT mapping. The alignment is a *structural correspondence*, not a formal equivalence. Each of pole structure, unitarity, and locality has a precise QFT formulation (the Källén–Lehmann representation; the S-matrix unitarity condition $S^\dagger S = 1$; cluster decomposition rather than naive locality alone), and the substrate-side conditions (C1)–(C3) align with the *structural content* of these QFT requirements without claiming term-by-term formal equivalence. Locality in QFT, in particular, is stronger than "records cluster in time and space" — cluster decomposition is the closer field-theoretic statement, and (C3) maps onto cluster decomposition rather than onto strict spacelike commutativity. The substantive content of the alignment is that the substrate-level conditions necessary for propagation pick out the same structural features that QFT independently identifies as necessary for a sensible field theory; the correspondence is at the level of these structural roles, not at the level of formal identity between (C1)–(C3) and any specific set of QFT axioms. A reader sceptical of the alignment can read it as: "the same kinds of conditions, motivated by different routes, end up doing the same work" rather than "(C1)=pole structure" as an equation.

At the closure threshold all three conditions fail simultaneously:

- (C1) fails because closure events define the limit of admissible distinguishability: at the boundary there are no further distinguishable successor states to support distinct poles.
- (C2) fails because the TPB rate saturates: there are no available ticks left to support unitary evolution at the propagation rate.
- (C3) fails because commitment events at the threshold are not coherent fragments of a record — they are the boundary conditions on what records can be, so they cannot themselves participate in local correlation functions.

We summarise:

Closure Non-Propagation Principle (CNPP). Closure-layer structure cannot appear as propagating quanta within the record layer, because the propagation conditions (C1)–(C3) — equivalently, the QFT requirements of pole structure, unitarity, and locality — fail at the closure threshold by construction. Closure-layer structure is therefore *constitutive* of the record layer, not *dynamical* within it.

The CNPP underwrites the OSP. Together they entail that the standard hierarchy-problem integral (4) is, when read naively to M_P, an extrapolation of a record-layer construction beyond the domain on which its integrand has physical content.

This is structurally analogous to demanding that a phonon dispersion relation remain physical above the Debye scale, where the lattice supporting the phonon dissolves. In that case the breakdown is uncontroversial because we already understand phonons as emergent. The VERSF claim is that the Higgs sector stands to the substrate as the phonon stands to the lattice.

The analogy has a limit worth naming. The Debye breakdown is geometrically visible: above the Debye scale, the wavelength of the would-be phonon falls below the lattice spacing, and the geometry alone tells you the description must fail. The VERSF breakdown is more abstract: TPB saturation, admissibility violation, and CCC failure are conditions on the substrate's bit-and-tick bookkeeping rather than on any visible geometry. Naming this dis-analogy strengthens the structural claim — VERSF's failure mode is of the same kind as Debye's (an emergent description ceasing to apply where its constituent substrate dissolves), but its mechanism is informational rather than spatial. The phonon case happens to be the one in which the mechanism is also geometrically transparent.

4. Admissibility Closure as a UV Suppression Kernel

The OSP and CNPP are negative results: they show what the standard argument cannot do. We now exhibit the positive mechanism. Admissibility closure regulates record-layer correlation functions by acting as a non-local UV suppression kernel.

A record-layer correlation function is computed not as an unconstrained path integral

$\langle\varphi(x)\varphi(y)\rangle_{naive} = \int \mathcal{D}\varphi\ \varphi(x)\varphi(y) \exp(iS[\varphi]/\hbar)$, (5)

but as an admissibility-restricted average

$\langle\varphi(x)\varphi(y)\rangle_{VERSF} = \int \mathcal{D}\varphi\ \delta(\mathcal{A}[\rho_\varphi])\ \varphi(x)\varphi(y) \exp(iS[\varphi]/\hbar)$. (6)

The projection $\delta(\mathcal{A}[\rho])$ restricts the integration to admissible configurations. In momentum space it manifests as a *suppression kernel* $F_{adm}(k)$ acting on each loop momentum:

$\int d^4k \to \int d^4k \cdot F_{adm}(k)$, (7)

with the defining property

$F_{adm}(k) \longrightarrow 0$ as $k \longrightarrow k_C$, (8)

where k_C is the relevant closure threshold. The kernel is not introduced by hand: its functional form is fixed by the admissibility constraint $\mathcal{A}[\rho]$ and the TPB-rate saturation, both of which follow from VERSF primitives.

4.1 The kernel and its substrate origin

The admissibility kernel $F_{adm}(k)$ is not introduced by hand. Its substrate origin is the κ-field — the propagating fluctuation of the commitment density around its committed background, $\rho(x,t) = \rho_0(x,t) + \kappa(x,t)$ — whose dynamics are derived from a Lagrangian in the VERSF programme (Taylor, *The Memory Kernel from First Principles*; Taylor, *On the Equivalence of the κ-Field and the Commitment Field*). The κ-field satisfies a damped massive Klein–Gordon equation with retarded Green's function

$G_\kappa(\tau) = \theta(\tau) \cdot e^{-\gamma_m \tau} \cdot \sin(\omega_\kappa \tau) / \omega_\kappa,\ \omega_\kappa = \sqrt{m^2_\kappa - \gamma^2_m}$, (9)

and a temporal memory kernel that takes the spatially-integrated asymptotic form

$K_{eff}(\tau) \sim A \cdot \cos(\omega_\kappa \tau + \varphi) / \tau,\ \tau \to \infty$. (9′)

The $1/\tau$ algebraic tail is a theorem of the retarded structure of a spatially extended source, not an assumed shape. The κ-field mass m_κ is itself derived (see §4.3 below).

The momentum-space suppression $F_{adm}(k)$ entering the regulated loop integrals (7) is the Fourier-side counterpart of this structure. Its substrate origin can now be made precise by appeal to a separate result. The κ-field response to a source depends on the source geometry: for a single point source the retarded Green's function decays as $\tau^{-3/2}$; for a distributed three-dimensional source the spatially-integrated Green's function is *exactly* $\sin(m_\kappa \tau)/m_\kappa$ (Taylor, *κ-Field Wave Dynamics, Geometric Memory, and Non-Markovian Decay*, §3); for a transversely narrow worldline tube it reduces to the $\cos(m_\kappa \tau)/\tau$ form of (9′). Quantum loop integrals — convolutions over a distributed vacuum field configuration rather than over a single worldline — fall in the first regime: the relevant κ-field response is the 3D-ensemble form $\sin(m_\kappa \tau)/m_\kappa$, whose Fourier transform is the Lorentzian

$$\tilde{G}_\kappa(p) = (p^2 + m^2_\kappa)^{-1}. \quad (10)$$

This identifies the *characteristic scale* of $F_{adm}(k)$ with m_κ as a theorem, not a structural identification: the κ-field's spectral response is centred on m_κ with a fixed Lorentzian profile, and $F_{adm}(k)$ inherits this scale exactly. The Lorentzian alone, however, is not by itself enough to produce a finite k_C^2 result for the regulated self-energy. A single propagator factor inserted into the schematic (11) below gives $\int k \cdot (k^2 + m^2_\kappa)^{-1}\, dk$, which still diverges logarithmically at large k. The Lorentzian thereby converts a quadratic divergence to a logarithmic one but does not finish the job. To produce the finite $(g^2/32\pi^2)\, k_C^2$ result reported in (11), the $F_{adm}(k)$ appearing under the integral must fall *faster* than $1/k^2$ at high k — supplied either by a doubled propagator structure, by an exponential cutoff, or by the full admissibility weighting that the master action is expected to deliver. This is the load-bearing remaining gap in §4: the Lorentzian fixes the *scale* exactly; the *suppression strength* needed to render the regulated integral finite still depends on a master-action $F_{adm}(k)$ that has not yet been derived. The illustrative forms

$$F_{adm}(k) \sim \exp(-(k/k_C)^n), \text{ with } n > 0, \quad (10')$$

are placeholders for this stronger suppression, not the Lorentzian itself; they encode the structural property (8) — $F_{adm}(k) \to 0$ as $k \to k_C$ — that the master-action kernel must satisfy and that drives (11). What is now anchored at the level of scale: the Lorentzian structure of (10) is derived exactly from the wave dynamics. What remains conditional at the level of functional form: the full $F_{adm}(k)$ that produces (11). The former is a theorem; the latter is the principal open derivation in this section.

What is and is not proved in §4.1. A clean separation:

Proved (Taylor, κ-Field Wave Dynamics paper): For distributed sources — the regime relevant to quantum loop integrals — the κ-field response is exactly $G_{eff}(\tau) = \sin(m_\kappa \tau)/m_\kappa$ in coordinate space, whose Fourier transform is the Lorentzian $\tilde{G}_\kappa(p) = (p^2 + m^2_\kappa)^{-1}$. This fixes the *characteristic scale* of $F_{adm}(k)$ at m_κ as a theorem, and is sufficient to convert the standard quadratic divergence into a logarithmic one in the schematic (11).

Outstanding (master-action derivation): The Lorentzian profile alone supplies only power-law suppression beyond m_κ and so leaves a logarithmic residue rather than producing the finite $(g^2/32\pi^2)\, k_C^2$ result that (11) reports. Closing this gap requires deriving the precise functional form of $F_{adm}(k)$ from the admissibility weighting in the VERSF master action — specifically, demonstrating that admissibility imposes a falloff stronger than $1/k^2$ at high k. The illustrative forms in (10′) — Gaussian, exponential — encode the required structural property but are placeholders, not derivations.

The size of the remaining work should not be misread. The regulation *mechanism* is in place (the substrate's admissibility weighting regulates loop integrals; this is the content of (7)–(8) and is proved). The *scale* of the regulation is in place (it is m_κ, by the Lorentzian). What is missing is the specific functional shape that turns log into finite — a single derivation, not a missing mechanism. The §4 regulator argument is therefore complete at the structural level and complete-plus-conditional at the quantitative level, with the conditionality precisely localised.

The one-loop scalar self-energy in the leading Yukawa channel reads, after Wick rotation to Euclidean momentum and standard angular integration over the 4-sphere,

$$\delta m_H^2 \sim (y^2 / 8\pi^2) \int_0^\infty dk_E \cdot k_E^3 \cdot G(k_E)^2 \cdot F_{adm}(k_E),$$

where $G(k_E) \sim 1/k_E^2$ is the massless internal propagator and the angular factor $2\pi^2$ has been absorbed into the prefactor. For a single propagator structure the $k_E^3 \cdot G(k_E)^2$ combination reduces to k_E, leaving the schematic

$$\delta m_H^2 \sim (y^2 / 16\pi^2) \int_0^\infty dk \cdot k \cdot F_{adm}(k) \sim (y^2 / 32\pi^2)\, k_C^2 \quad (11)$$

for any kernel satisfying (8). The Λ^2 coefficient familiar from cutoff regularisation is recovered in the limit $F_{adm} \to 1$ with $k_C \to \Lambda$; the VERSF kernel replaces this hard cutoff with a substrate-derived smooth suppression and yields a finite, kernel-controlled result.

Two qualifications are in order. First, (11) captures only the leading Yukawa contribution; the full one-loop scalar self-energy receives gauge-boson and Higgs-self-coupling contributions of structurally identical form, each with its own coupling and channel-dependent O(1) prefactor. The dominant numerical factor remains k_C^2, with channel-summed coefficient of order $(g^2/16\pi^2)$ where g denotes the relevant coupling. Second, the integral as written assumes a massless internal propagator; including m^2 in the denominator modifies the IR behaviour but not the UV scaling that drives (11).

If k_C were identified with M_P, this would reproduce the standard catastrophe. But under the OSP and CNPP, k_C is *not* M_P. It is the effective closure threshold of the record layer, set by substrate-level dynamics — specifically, by the κ-field mass derived in §4.3. The remainder of §4 develops this identification.

The mass correction (11) is therefore *finite* and parametrically of order $(y/4\pi)^2\, k_C^2$, not $(y/4\pi)^2\, M_P^2$. The seventeen-order hierarchy is replaced by a single-order separation between v and k_C, set by the width of the admissibility crossover.

Epistemic status. That admissibility closure regulates record-layer correlation functions via a suppression kernel of the form (7)–(8) is *proved* within VERSF (Admissibility Closure paper). The κ-field memory kernel $K(\tau)$ is *derived* from the κ-field Lagrangian — taking the form $\sin(m_\kappa \tau)/m_\kappa$ for distributed sources (the regime relevant to loop integrals) and the asymptotic form $A \cos(\omega_\kappa \tau + \varphi)/\tau$ for the worldline-tube regime; the Fourier transform of the distributed-source kernel gives the Lorentzian (10) with characteristic scale exactly m_κ. The illustrative form (10′) is retained as one concrete realisation of the structural property (8) that drives (11). That k_C is identified with the κ-field mass m_κ, itself derived from the K=7 architecture and the CCC threshold, is established in §4.3 below.

4.2 The emergent record-layer closure scale

The relationship between M_P and k_C deserves direct comment. A natural objection runs: if the substrate closure threshold is M_P, why should k_C — the effective threshold above which record-layer coherence breaks down — sit many orders of magnitude lower?

The answer is that coherent record-layer transport requires more than mere substrate admissibility. It requires that commitment events compose into stable, causally consistent, distinguishable records. This is the content of the *Causal–Coherence Compatibility* (CCC) condition formalised in Taylor, *Causal–Coherence Compatibility and the Fact-Production Threshold*. In brief: irreversible fact formation requires a minimum action budget $\rho L^4 \gtrsim \hbar c$ within a region of size L, defining the *coherence scale*

$$\xi = (\hbar c / \rho)^{1/4} \quad (12)$$

as the smallest spatial extent capable of sustaining a committed fact. CCC is the stronger requirement that a *sequence* of commitment events compose into a record: (i) successive events stand in admissible causal relation, *and* (ii) the resulting record remains coherent — its distinguishability structure is preserved across the sequence — *and* (iii) the joint condition (i)+(ii) holds simultaneously rather than at the expense of one another. Configurations can satisfy admissibility yet fail CCC because the causal and coherence demands pull against each other at high commitment rates. The CCC threshold is therefore a strictly stronger condition than admissibility alone, and it operates at the scale ξ rather than at ℓ_P.

As one approaches M_P from below, CCC fails well before substrate admissibility does. The effective record-layer closure scale k_C may therefore be identified with the energy scale at which CCC fails for sustained record transport, even though substrate admissibility remains formally well-defined up to M_P itself. §4.3 makes this identification quantitative by deriving k_C from the κ-field mass.

This pattern is not unique to VERSF. It is characteristic of emergent coherence systems:

- **Hydrodynamics** breaks down at scales well above the atomic scale, where mean-free-path coherence fails — not at the atomic scale itself, where the underlying particles still exist.

- **Phonons** decohere at scales well below lattice collapse, where anharmonic and impurity effects spoil their quasiparticle status — not at the chemical-bond scale, where the lattice still nominally exists.
- **Superconducting coherence** is lost at temperatures well below the electronic-binding scale, when Cooper-pair stability fails — not at the scale at which electrons themselves cease to exist.

In each case the *coherence breakdown* scale sits orders of magnitude below the *constituent breakdown* scale. The two scales are categorically distinct, and their separation is not fine tuning but a generic feature of layered emergent systems.

The VERSF claim is that the hierarchy between v and M_P fits this pattern. The seventeen-order separation is not a precise cancellation between bare and counterterm masses; it is a coherence-versus-constituent separation of exactly the kind one finds in every well-understood emergent system. What needs to be derived is not the existence of the separation but its specific magnitude, which is set by substrate parameters governing CCC and TPB.

4.3 The closure scale from the κ-field mass

What is and is not derived. The *form* $k_C = \sqrt{(4/3)} \cdot \xi^{-1}$ is locked by the K=7 spectral data and the CCC threshold (§§4.3.1–4.3.2 below). The *value* of k_C, however, depends on which ξ governs the electroweak sector, and that ξ has not yet been derived from the master action. The cosmological CCC scale $\xi \approx 8.2 \times 10^{-5}$ m (Taylor, *A Hidden "Middle Scale" in the Universe*) gives $\xi^{-1} \approx 2.4$ meV and $m_\kappa = \sqrt{(4/3)} \cdot \xi^{-1} \approx 2.8$ meV — fourteen orders of magnitude below the electroweak scale. With $k_C \equiv m_\kappa$ at that value, the regulated mass correction (11) becomes $\delta m_H^2 \sim (g^2/32\pi^2) \cdot m^2_\kappa \sim 10^{-6}\ \mathrm{eV}^2$, which is not the observed Higgs scale and is not even the right order of magnitude for any electroweak phenomenon. So either (a) the relevant ξ for the EW sector is sector-specific and much shorter than the cosmological CCC scale, or (b) the §4 regulator argument as written is not the mechanism doing the EW work, and a different closure-scale identification is needed. The §4.3.3 "open identification problem" subsection below lays out three structural responses; the honest reading is that the §4 regulator argument supplies the *form* of the answer but not yet the *number*. The framework's structural claims — the category-error thesis (§§2–3), the OSP and CNPP (§§2.1, 3), the regulator-dependent reading of the quadratic divergence (§5), and the coherence-eigenmode reading of the Higgs (§6) — do not depend on this identification and are unaffected by it. What does depend on it is the quantitative claim that the §4 mechanism *numerically* regulates the Higgs self-energy at the observed scale.

The qualitative claim that coherence breakdown precedes constituent breakdown is generic. The quantitative claim that $k_C \ll M_P$ follows from a derivation of the κ-field mass m_κ, with which k_C is identified at the structural level. The remaining quantitative work is the identification of the relevant ξ for the EW sector.

The derivation

The κ-field κ(x,t) is the propagating fluctuation of the commitment density ρ(x,t) around its committed background $\rho_0(x,t)$. Its dynamics are governed by a Klein–Gordon Lagrangian whose mass term is *not* a free parameter but the Hessian of the commitment free-energy functional $F(\rho)$ at the equilibrium background (Taylor, *The Memory Kernel from First Principles*; Taylor, *The κ-Field Mass as the Physical Hessian of the Commitment Constraint Surface*):

$m^2_\kappa = F''(\rho_0)$. (13)

Under the K=7 minimal fact architecture (Taylor, *A No-Go Theorem for Non-Simplicial Relational Substrates – K=7*) and the CCC coherence-scale relation (12), this evaluates to

$m^2_\kappa = \lambda_{eff} \cdot \xi^{-2} = (4/3) \cdot \xi^{-2}$, (14)

where $\lambda_{eff} = 4/3$ is the lowest positive eigenvalue of the closure operator $L_{eff} = (4/3)\, I_4$ on the physical state space V_p, derived from the spectral structure of the Fano-plane incidence matrix of PG(2,2). The derivation is locked at every step: the Uniqueness Theorem for the constraint penalty (Schur's Lemma applied to the irreducible PGL(3,2) representation on V_p), the No-Alternative-Scaling Lemma (CCC sector reduces all dimensional quantities to ξ alone), and the Constraint-to-Mass Proposition (V_c can only enter an admissible Lorentz-invariant Lagrangian as a Klein–Gordon mass term) together force (14) as a theorem rather than an ansatz. No free parameters remain.

The κ-field is the universal mediator of record-layer dynamics: it propagates the commitment-density response of every prior commitment event forward through time. The natural closure scale for the record layer is therefore the scale above which κ-field coherence cannot be sustained — and that scale is m_κ itself. We identify

$k_C \equiv m_\kappa = \sqrt{(4/3)} \cdot \xi^{-1}$. (15)

The hierarchy $k_C \ll M_P$ then becomes

$k_C / M_P = \sqrt{(4/3)} \cdot \ell_P / \xi$, (16)

a ratio set entirely by the small parameter ℓ_P / ξ — the Planck length divided by the CCC coherence scale. The seventeen-order gap between v and M_P is reduced to a single derived dimensionless ratio, evaluated at substrate parameters fixed by the K=7 spectral data and the CCC threshold.

The informational reading

The mass derivation has a transparent informational interpretation that complements the formal result. A coherence mode is sustained by an admissibility budget bounded by area-law (Bekenstein-style) considerations: a region of size L hosts at most $\sim (L/\ell_P)^2$ distinguishable admissible configurations, not the volume-law $(L/\ell_P)^3$ that naive cell-counting would suggest. An SM-style coherence mode requires a substantial bit content N_{bits} — encoding gauge-group structure, phase resolution, and correlation across the coherence volume — and the smallest L_{coh} capable of hosting such a mode satisfies

$(L_{coh} / \ell_P)^2 \gtrsim N_{bits} \Longrightarrow \omega_{max} \sim M_P / \sqrt{N_{bits}}$.

This square-root-in-N_bits scaling is compatible with the parallel character of coherence (the substrate's ticking parallelises across the volume, so a naive serial tick-counting argument such as $\omega_{max} \sim M_P / N_{bits}$ fails) and yields a suppression below M_P set by the area-law coordination cost. For an SM-style coherence mode N_{bits} is large but not yet derived from the master action; the bit-budget argument gives the *form* of the suppression — $k_C \sim M_P / \sqrt{N_{bits}}$ with $\sqrt{N_{bits}} \gg 1$ — but does not yet supply a numerical value. The corresponding identification $\xi^{-1} \sim M_P / \sqrt{N_{bits}}$ is equivalent to $L_{coh} \sim \xi$ in the area-law accounting. The κ-field mass *is* the field-theoretic realisation of the area-law coordination cost; the informational and field-theoretic descriptions are dual views of the same structure, with neither view yet supplying a numerical k_C for the EW sector.

The remaining open question

What remains open is not the *form* of the derivation but the *identification of the relevant ξ* for the electroweak hierarchy. The cosmological CCC coherence scale derived in Taylor (*A Hidden "Middle Scale" in the Universe*) from vacuum-energy considerations is $\xi \approx 8.2 \times 10^{-5}$ m, giving $\xi^{-1} = \hbar c/\xi \approx 2.4$ meV and $m_\kappa = \sqrt{4/3} \cdot \xi^{-1} \approx 2.8$ meV in natural units — fourteen orders of magnitude below the electroweak scale.

The leading reading: the geometric-mean rung (option (iii)). The paper's quantitative claim for the EW sector is that the §4 regulator mechanism applies at the §8 geometric-mean scale, not at m_κ. That is, the closure-scale identification (15) supplies the *form* of the answer — $k_C = \sqrt{4/3} \cdot \xi^{-1}$ — but for the EW sector specifically, the regulator-effective scale is the emergent rung E_{geo} on the closure–CCC ladder. Substituting $m_\kappa \approx 2.78$ meV $\approx 2.78 \times 10^{-12}$ GeV into the geometric mean developed in §8 gives

$$E_{geo} = \sqrt{M_P \cdot m_\kappa} \approx \sqrt{1.22 \times 10^{19} \text{ GeV} \times 2.78 \times 10^{-12} \text{ GeV}} \approx 5.8 \text{ TeV},$$

within one order of magnitude of $v = 246$ GeV and inside the §8 coherence band by construction. The schematic loop integral (11) holds under this reading with k_C replaced by E_{geo}: the §4 mechanism supplies the regulation form, the §8 ladder structure supplies the operational scale, and the two are not parallel applications of the same mechanism but a single combined reading. The §4 and §8 arguments are unified under option (iii) as: the §4 regulator mechanism applied at the §8 geometric-mean scale.

Alternative resolutions consistent with the structural form. Options (i) and (ii) below would obtain if a master-action derivation later established a different ξ for the EW sector. They are not currently competing predictions with (iii) — they are placeholders for "some other ξ exists, to be derived" — but the structural form $k_C = \sqrt{4/3} \cdot \xi^{-1}$ would survive intact under either:

(i) the relevant ξ for the EW sector is a *distinct* coherence scale of the same architecture, set by sector-specific substrate parameters rather than by the cosmological vacuum-energy density. Under this reading the identification $k_C \equiv m_\kappa$ of (15) applies directly with a sector-specific

m_κ in the EW range, and the §4 regulator mechanism does its work at k_C in the conventional reading.

(ii) the hierarchy involves *multiple* closure scales corresponding to different layers or sectors of the κ-field, with the EW-relevant k_C selected by sector-specific commitment dynamics. Like (i), this preserves $k_C \equiv m_\kappa$ at a sector-appropriate scale.

Status of the identification (15) across the readings. Under (iii) — the leading reading — the regulator-effective scale for the EW sector is $E_{geo} \approx 5$–6 TeV, *not* $m_\kappa \approx 2.8$ meV; m_κ enters as the IR input to E_{geo} rather than as the regulator threshold itself. Under (i)/(ii) — the alternative resolutions — the identification $k_C \equiv m_\kappa$ applies directly, with the EW-relevant m_κ supplied by a sector-specific ξ awaiting derivation. The structural form (15) is preserved across all three; what differs is the operational scale.

Distinguishing (iii) from (i)/(ii) — and ultimately deriving the EW-relevant scale from the VERSF master action — is the principal remaining task. What is established is the *form* of the answer: the closure-scale identification (15) is locked by the K=7 architecture and the CCC threshold, the seventeen-order gap between v and M_P is the difference between two derived structural scales rather than a fine-tuning cancellation, and the leading reading (iii) places the EW scale within an order of magnitude of where it is observed using inputs fixed by independent results. What is not yet established is what selects v within the coherence band that (iii) identifies (likely a TPB-rate ratio from sector-specific commitment dynamics).

---

5. Quadratic Divergences and Regulator-Dependent Interpretation

Role in the overall argument. This section makes a supporting observation, not an independent argument. The load-bearing reply to the conventional hierarchy framing is the OSP/CNPP claim of §§2–3: the Planck scale is not a sector of states from which threshold corrections to m_H^2 could arise, because it is a closure threshold of the substrate rather than a UV completion of the record layer. The present section adds that *even granting* the standard cutoff-regularisation framing, the catastrophic Λ^2 coefficient has no scheme-independent physical content — it is a feature of one regularisation scheme that the standard naturalness narrative reads as if it were physics. This is an auxiliary point: it weakens the rhetorical force of the Λ^2 catastrophe without itself replacing OSP/CNPP as the structural answer. The hierarchy problem's deep content — *if heavy states exist coupled to the Higgs, their threshold corrections to m_H^2 are of order $M^2_{heavy}/(16\pi^2)$, regardless of regulator* — is answered by OSP/CNPP (no such heavy states at M_P), not by regulator choice. The argument below should be read as supplementary to §§2–3, not as a stand-alone leg.

In dimensional regularisation, where loop integrals are performed in $d = 4 - 2\varepsilon$ dimensions, scalar self-energies do not develop power-law divergences at all. The would-be quadratic sensitivity is absent; only logarithmic divergences of physical masses appear. The "Λ^2

catastrophe" is therefore a feature of cutoff regularisation — in which a dimensionful scale Λ is inserted by hand to truncate momentum integrals — rather than a scheme-independent prediction.

This is well known to specialists but is often elided in the standard naturalness narrative. The honest position is that:

- Cutoff regularisation produces a Λ^2 sensitivity that lacks an unambiguous physical interpretation.
- Dimensional regularisation produces no such sensitivity.
- The physical content — the *finite* radiative corrections to m_H — is regulator-independent and modest.

The naturalness argument then depends on a specific reading of cutoff regularisation in which Λ is endowed with physical meaning as the scale of new degrees of freedom. Without (P2) — without identifying Λ with a record-layer UV completion — the Λ^2 coefficient has no physical referent.

VERSF turns this observation into a structurally stronger claim than is available in the conventional dimensional-regularisation argument. Bardeen's well-known position — that only logarithmic divergences carry scheme-independent physical content, and that the quadratic sensitivity is a feature of one particular regularisation scheme — relies on an argument that one renormalisation scheme is, in some sense, "more physical" than others. This is a defensible position but a meta-theoretical one: it adjudicates between schemes from outside the schemes themselves.

VERSF avoids this manoeuvre entirely. The admissibility kernel of §4 is not a choice of scheme; it is a physical regulator derived from substrate dynamics. The question "which scheme is right?" is replaced by the question "which physical mechanism does the regularising?" — and the answer is admissibility closure, which is independently motivated by the VERSF primitives. The resulting finite mass correction (11) is regulator-independent in the sense that any kernel satisfying (8) yields a correction of order $(g^2/16\pi^2)\, k_C^2$. The apparent severity of the hierarchy problem in cutoff-regularisation language is therefore, in part, a consequence of reading regulator-dependent scheme behaviour as if it were scheme-independent physics — a confusion the VERSF framework dissolves by supplying an unambiguous physical regulator from substrate primitives rather than from a scheme preference.

The Λ^2 in equation (1) is not a measurable quantity. It is a feature of one regularisation scheme that happens to be in widespread expository use. The VERSF kernel, by contrast, derives from the substrate's admissibility structure and produces finite, scheme-independent corrections.

This does not by itself dissolve the hierarchy problem — one must still explain why the EW scale sits at 246 GeV rather than elsewhere — but it changes what *kind* of explanation is required. It is no longer a problem of cancellation; it is a problem of substrate dynamics. The structural answer to the cancellation question is OSP/CNPP, as developed in §§2–3; the present section's role is to remove the rhetorical scaffolding (the catastrophe-language of "10^{32} fine-tuning") that makes the cancellation framing seem inevitable.

---

6. The Higgs as a Record-Field Coherence Eigenmode

Role in the overall argument. Sections 2–5 and §10 do the load-bearing work for the category-error claim: the OSP and CNPP (§§2–3) establish that the Planck scale is not a sector of states from which the Higgs receives threshold corrections, the §4 admissibility kernel supplies a physical regulator at substrate-derived scales, and §5 removes the regulator-dependent rhetorical scaffolding that makes the conventional Λ^2 framing seem inevitable. The present section addresses the third premise of the conventional argument — that the Higgs is a fundamental scalar — but does not by itself dislodge it. The coherence-eigenmode reading sketched here is a *constructive target*: a research direction that the structural argument creates, sitting alongside the master-action derivation of $F_{adm}(k)$ (the §4.1 outstanding work) as the second item on the programme's near-term to-do list. The functional (17) below has five coefficients none of which are derived from the master action; the vacuum and the 246 GeV scale are not produced. What this section delivers is a *sketch of what dislodging (P3) would look like* within VERSF — not a derivation. A reader following the load-bearing argument can take §6 as a directional commitment; the category-error claim does not depend on the constructive programme of §6 being already complete.

A note on language across the VERSF programme. The "record field" introduced below is the field-theoretic manifestation of the deeper scalar commitment / entropic structure developed across the wider VERSF programme. Earlier papers in the programme often described this structure in entropy-centric language — irreversible commitment, distinguishability bookkeeping, closure-class counting — because the emphasis there was on the substrate's own dynamics. The present paper adopts the language of coherence modes and emergent field structure because the focus is the Higgs sector and its relation to effective field theory. The underlying ontology is the same in both cases: stable field-theoretic excitations on the record layer are organised commitment-density configurations on the substrate, and the commitment density ρ used here is the same quantity whose closure structure the entropy papers analyse. The terminological shift is one of register, not of object — a reader who has encountered VERSF through the entropy or κ-field papers should read "record field" here as the field-theoretic face of structures already familiar in that other language.

Sections 2–5 dispose of premise (P2) and weaken premise (P1). We now sketch a constructive approach to premise (P3): the assumption that the Higgs is a fundamental scalar.

In VERSF, the record layer supports stable coherence eigenmodes of the record field, built on the commitment density ρ. The minimal phenomenological description used in Section 2 (a Ginzburg–Landau scalar) is too generic; the VERSF-specific functional must include record-current transport and admissibility coupling explicitly:

$$\mathcal{C}[\rho] = \int d^4x \, [\, \alpha\, |\nabla\rho|^2 + \beta\, \rho^2 + \gamma\, \rho^4 + \eta\, (\nabla\cdot J_\rho)^2 + \lambda\, \mathcal{A}[\rho]^2 \,], \quad (17)$$

where

- $\rho(x)$ is the local commitment density,
- $J_\rho(x)$ is the record-current encoding coherent commitment transport,
- $\mathcal{A}[\rho]$ is the admissibility-closure functional,
- $\alpha, \beta, \gamma, \eta, \lambda$ are not free parameters but functionals of substrate parameters (TPB rate, BCB normalisation, closure thresholds).

The $(\nabla\cdot J_\rho)^2$ term penalises deviations from local conservation of the record-current J_ρ: configurations in which the current fails to be locally conserved cost coherence energy proportionally to the square of the divergence. It is motivated by the BCB requirement that commitment events balance globally — local violations of the corresponding current's conservation register as coherence-energy penalties rather than being strictly forbidden. The $\mathcal{A}[\rho]^2$ term plays the analogous role for admissibility: configurations violating admissibility cost coherence energy proportionally to the square of their violation. It is the natural soft-constraint counterpart to the hard projection $\delta(\mathcal{A}[\rho])$ introduced in §4.

These two ingredients are *structurally distinctive* to VERSF rather than uniquely VERSF in any decisive sense — strictly comparable terms could in principle be written in any framework that supplies analogous primitives. What is distinctive is the *motivation*: the $(\nabla\cdot J_\rho)^2$ and $\mathcal{A}[\rho]^2$ terms here are not added by hand to fit phenomenology but follow from BCB and admissibility-closure structures that are independently present in the VERSF programme. The precise coefficients η and λ, however, await derivation from the master action.

Stationary configurations satisfy

$\delta\mathcal{C}/\delta\rho = 0$, (18)

and the symmetry-broken vacuum at $\langle\rho\rangle = \rho_0$ corresponds to a coherent record-field ground state. Fluctuations about ρ_0 furnish the analogue of the Higgs scalar; its mass is set by the curvature of $\mathcal{C}$ at ρ_0,

$m_h^2 = (\partial^2\mathcal{C}/\partial\rho^2)|_{\rho_0} / Z(\rho_0)$, (19)

with Z a wave-function normalisation. The mass is not a bare parameter; it is a derived eigenvalue of the substrate-determined functional.

This is the same logical structure as the BCS gap, the σ-mode of chiral symmetry breaking, or the superfluid order parameter — but with two distinctive features. First, the underlying substrate is *pregeometric*, not a field-theoretic gauge sector. Second, the admissibility constraint is explicit in (17), supplying the suppression mechanism of Section 4 from within the same functional.

Epistemic status. The functional form (17) is *conditional*: the inclusion of $(\nabla\cdot J_\rho)^2$ and $\mathcal{A}[\rho]^2$ terms follows from VERSF primitives, but the precise coefficients require derivation from the master action. That fluctuations about $\langle\rho\rangle$ reproduce the SM Higgs in full (couplings to W, Z, fermions; observed branching ratios) is *conjectural* and constitutes a concrete research target.

7. Why This Is Not Composite Higgs

The emergence reading invites an immediate conflation. Composite-Higgs models also treat the Higgs as a bound state of more fundamental dynamics. The VERSF reading must be distinguished sharply.

Composite Higgs. A new strongly-coupled gauge sector at the scale $\Lambda_{TC} \sim 10$ TeV produces a bound-state pseudo-Nambu-Goldstone boson identified with the Higgs. The "constituents" are *fields* of the same ontological kind as the Standard Model fields: they live in the record layer, propagate, carry charges, and admit a Lagrangian description in their own right. The emergence is *within field theory*. Naturalness is restored at Λ_{TC}, beyond which the standard problem reappears in a different sector.

VERSF. The Higgs is a coherence eigenmode of the record field, built on the commitment density ρ. The "constituents" are commitment events of the substrate — *not* fields, not propagating quanta, not gauge-charged. They are pre-field elements of the closure layer, which does not admit a Lagrangian description because the conditions for propagation (CNPP) are absent. The emergence is *from a pre-field closure substrate*. There is no analogue of Λ_{TC} at which the naturalness question recurs: the closure threshold is not a UV completion in the field-theoretic sense.

The non-recurrence is the categorical distinction. Composite Higgs faces a recurrence: as LHC null results push Λ_{TC} upward, Λ_{TC} itself becomes increasingly far from the EW scale and itself needs protection against still-higher physics. The naturalness question reappears one level up, now between Λ_{TC} and whatever sits above it. This is generic to any field-theoretic UV completion: every new scale is, in turn, a candidate for the same hierarchy problem one level higher. VERSF predicts *no such recurrence* because the closure threshold is not a UV completion. It is not a sector of states that could itself receive radiative corrections from a higher sector — it is the limit of admissible distinguishability, the boundary at which "record-layer state" is no longer the right ontological category. The category-error reading dissolves the hierarchy problem at the EW–Planck pair; the absence of recurrence is what makes the dissolution categorical rather than just one more shift of the same problem.

The distinction can be put schematically:

Composite Higgs: Higgs = bound state of *fields* at $\Lambda_{TC} \subset \mathcal{L}_R$. VERSF: Higgs = coherence eigenmode of *commitments* on $\mathcal{L}_C$, observed in $\mathcal{L}_R$.

Composite Higgs replaces one record-layer Lagrangian with another. VERSF replaces the record-layer description, at its UV boundary, with a categorically different substrate. Composite Higgs is a horizontal move within field theory; VERSF is a vertical move out of it.

A practical consequence: composite Higgs predicts a tower of techni-resonances near Λ_{TC}, accessible at the LHC and FCC. VERSF predicts no such tower — only the smooth turn-off of the suppression kernel $F_{adm}(k)$ near k_C, and the absence of clean particle states above it.

---

8. The Geometric-Mean Relation as a Coherence Band

This section develops option (iii) of §4.3: the geometric-mean ladder rung. It is no longer presented as an independent consideration that merely corroborates the layered reading. Under the v8 reconciliation, §8 *is* the option (iii) calculation — the only one of the three §4.3 responses that currently makes numerical contact with the electroweak scale from inputs fixed entirely by independent results.

The structural claim. The κ-field mass m_κ derived in §4.3 from the K=7 architecture and the CCC threshold (applied to the cosmological vacuum-energy density), and the Planck mass M_P, are two substrate-derived scales. Their geometric mean

$E_{geo} = \sqrt{(M_P \cdot m_\kappa)} \approx \sqrt{(1.22 \times 10^{19}\ \text{GeV} \times 2.78 \times 10^{-12}\ \text{GeV})} \approx 5.8\ \text{TeV}$ (21)

places v = 246 GeV within one order of magnitude of where the framework predicts an intermediate coherence rung — and at a scale where the LHC has, to date, observed no new physics. This is not a derivation of 246 GeV. It is, however, a contact-with-data result the framework neither needed nor would have obtained had the K=7+CCC architecture been wrong. Every input on the right-hand side of (21) is fixed by considerations independent of the hierarchy problem itself: M_P from Newton's constant and the Planck units, m_κ from the K=7 spectral coefficient 4/3 combined with the observed cosmological vacuum-energy density via CCC. There are no parameters tuned to produce the EW result; the relation expresses a one-parameter coupling between three observed scales (ρ_Λ, M_P, v) through a structural coefficient derived from the substrate architecture. This is a one-parameter prediction between observables, not a derivation from first principles — but it is also not numerology, since the coefficient $\sqrt{(4/3)}$ is a theorem of the K=7 spectral data and is not adjustable.

Relation to prior observations of the same coincidence. The numerical observation $\sqrt{(M_P \cdot M_\Lambda)} \sim$ TeV has been noted in several earlier contexts. Dirac's large-numbers hypothesis (Dirac 1937, 1938) proposed structural relations between cosmological-scale and microphysical quantities; Banks and others have discussed cosmological-constant scaling within holographic and supersymmetry-breaking frameworks (Banks 2000); and the cosmological-constant-as-IR-cutoff literature (Cohen, Kaplan, and Nelson 1999) derived a similar geometric-mean scale from black-hole-entropy bounds on effective field theories. What is new in the VERSF reading is therefore not the coincidence itself, which has been recognised before, but its structural origin. The IR scale m_κ entering (21) is the κ-field mass derived from K=7 spectral data and the CCC threshold; the $\sqrt{(4/3)}$ coefficient that distinguishes m_κ from the bare M_Λ is the same K=7 spectral structure that gives $n_s \approx 1 - 2/N\star \approx 0.964$, compared with Planck 2018's $n_s = 0.965 \pm 0.004$. Two empirical contacts — a TeV-scale coherence band edge and a percent-level CMB scalar spectral index — issue from one substrate architecture. The geometric mean is therefore not a numerological observation but a structural consequence of the closure–CCC ladder. Earlier identifications of the coincidence pointed at a coincidence; the VERSF reading proposes a substrate-level reason.

The unification with §4.3. The IR scale entering (21) is the same one entering the §4.3 κ-field-mass calculation: a single coherence scale supplied by the CCC threshold $\rho\ \xi^4 = \hbar c$ applied to the cosmological vacuum-energy density. The §4.3 calculation reads it through the κ-field mass $m_\kappa = \sqrt{(4/3)} \cdot \xi^{-1}$; the §8 calculation reads it through the IR partner of a closure–CCC geometric mean. These are two readings of the same substrate parameter, not two independent applications.

The geometric-mean form. In length units, the geometric mean of the Planck length ℓ_P and the cosmological CCC coherence length $\xi_{CCC} \approx 8.2 \times 10^{-5}$ m is

$$\xi_{geo} = (\ell_P \cdot \xi_{CCC})^{1/2} \approx \sqrt{(1.6 \times 10^{-35}\ \text{m} \times 8.2 \times 10^{-5}\ \text{m})} \approx 3.6 \times 10^{-20}\ \text{m}, \quad (20)$$

corresponding via $E = \hbar c/\xi_{geo}$ to ≈ 5.4 TeV in energy units. The slight difference from (21)'s 5.8 TeV is just the $\sqrt{(4/3)}$ factor between $M_\Lambda = \xi^{-1}$ and $m_\kappa = \sqrt{(4/3)}\cdot\xi^{-1}$ — the energy-side geometric mean using M_Λ as the IR scale differs from the version using m_κ by $(4/3)^{1/4} \approx 1.075$, and 5.4 × 1.075 ≈ 5.8 TeV.[^geomean] Earlier drafts of this section wrote (20) with the Hubble horizon L_H in place of ξ_{CCC}; that form invokes a different IR scale ($1/L_H \sim H_0 \sim 10^{-33}$ eV) and yields a meV-scale geometric mean rather than a TeV-scale one. The correct IR scale for the present argument is ξ_{CCC}, fixed by the CCC threshold applied to the cosmological vacuum-energy density, since that is the scale entering m_κ in (15). The "coherence band" generated by (20)–(21) is therefore the substrate-derived intermediate band between Planck-scale closure structure and CCC-scale vacuum coherence, not between Planck-scale structure and horizon-scale cosmology.

[^geomean]: The two energy-side geometric means are $\sqrt{(M_P \cdot M_\Lambda)}$ and $\sqrt{(M_P \cdot m_\kappa)}$. Their ratio is $\sqrt{(m_\kappa/M_\Lambda)} = \sqrt{(\sqrt{(4/3)})} = (4/3)^{1/4} \approx 1.075$. Either is a legitimate one-parameter prediction; the §8 numerical result is robust within this factor.

Why the geometric mean is not arbitrary. Geometric-mean intermediate scales recur across physics whenever a dynamical system is bounded by two cutoffs with approximately scale-free *bracketing geometry* between them. The bracketing geometry is scale-free in the sense that no preferred intermediate scale is imposed *on* the interval from outside — the bracket itself has no built-in landmark — even if intermediate dynamical scales emerge *from* the bracket as predicted rungs of the ladder. The distinction is crucial: VERSF can predict a ladder of intermediate rungs (as §11.2 sketches) precisely because the bracket is scale-free at the geometric level and the rungs are derived as emergent dynamical scales rather than imposed as additional input. The pattern is not a coincidence; it is a structural feature of the bracketing geometry. Four observations make this concrete.

First, scale-free intermediate dynamics force geometric-mean output scales. When the dynamics between a UV scale Λ_{UV} and an IR scale Λ_{IR} are scale-invariant or near-scale-invariant — no additional intrinsic energy scale, only the two cutoffs — and the emergent dynamical scale depends on the inputs through a single power-law combination, dimensional analysis combined with reflection symmetry under $\Lambda_{UV} \leftrightarrow \Lambda_{IR}$ forces that combination into the form $\Lambda_{em} \sim (\Lambda_{UV} \cdot \Lambda_{IR})^{\alpha}$ with $\alpha = 1/2$ the unique symmetric power. The reflection symmetry itself is not arbitrary: in the absence of any structural distinction between the UV and IR cutoffs at the level of the bracketing geometry — no preferred direction, no asymmetric coupling — no mechanism is available to favour one over the other, so the symmetric power $\alpha = 1/2$ is selected. Other powers (1/3, 2/3, ...) are dimensionally consistent but require additional structural input distinguishing UV from IR, which the bracketing geometry alone does not supply. The single-power-law restriction is itself a consequence of scale invariance: a non-power functional dependence $f(\Lambda_{UV}, \Lambda_{IR})$ introduces an implicit additional scale and so violates the scale-free assumption. Within those assumptions, the geometric mean is therefore not chosen; it is *forced* by the absence of intermediate scales.

Second, holographic UV/IR coupling repeatedly produces this structure. The Cohen–Kaplan–Nelson bound (Cohen, Kaplan, and Nelson 1999) derives a TeV-scale UV cutoff $\Lambda_{UV} \lesssim (M_P^2 / \Lambda_{IR})^{1/2} \sim \sqrt{M_P \cdot M_\Lambda} \sim$ TeV from the requirement that the entropy of an effective field theory within an IR volume L^3_{IR} not exceed the entropy of a black hole of the same size. The bound is structural — it follows from black-hole thermodynamics, not from supersymmetry, conformality, or any other additional input — and it produces the same geometric-mean scale that (21) does. Banks (2000) reaches a similar conclusion within holographic and SUSY-breaking frameworks. The cosmological-constant-as-IR-cutoff literature more broadly identifies this geometric mean as a generic feature of theories in which gravitational and infrared scales are linked. VERSF inherits this pattern: it does not invent the geometric-mean form; it supplies the substrate-level mechanism that explains why the bracketing structure obtains.

Third, VERSF constrains the IR scale via CCC. What the prior literature leaves open — *which* IR scale enters the geometric mean, and *why that specific scale* — is what the VERSF reading fixes. The Causal–Coherence Compatibility threshold relates the cosmological vacuum-energy density to a single coherence scale, which in turn fixes the κ-field mass m_κ entering (21). The crucial point is that this scale is not a parameter chosen to make the geometric mean land at the right place: it is the IR scale already supplied by the substrate dynamics that, in a separate sector, gives $n_s \approx 0.964$ against Planck's 0.965 ± 0.004. The same scale plays both roles because the same architecture controls both.

Fourth, the TeV-scale coherence rung is therefore structurally expected. Given (i) the geometric-mean form forced by scale-free intermediate dynamics, (ii) the holographic precedent identifying $\sqrt{M_P \cdot M_\Lambda}$ as the relevant scale, (iii) the CCC mechanism fixing $M_\Lambda \sim m_\kappa$ from the cosmological vacuum-energy density, and (iv) the K=7 spectral coefficient $\sqrt{4/3}$ entering m_κ — the appearance of a coherence rung in the few-TeV range is not a numerological surprise. It is what one expects when these four structural elements combine. The empirical content is not "there exists a TeV-scale coherence rung" — that follows. The empirical content is the *consistency* of three independently measured scales (ρ_Λ, M_P, v) under the relation (21), and the further consistency of the LHC's null results with the smooth-turn-off phenomenology that (21) implies.

This is what is meant when (21) is described as a one-parameter prediction between observables rather than numerology. The geometric mean is structurally forced, the input scales are independently fixed, and the coefficient is theorematic. What remains adjustable — and is honestly flagged as such — is the further question of what selects v *within* the coherence band (likely a TPB-rate ratio derived from sector-specific commitment dynamics), not whether the band sits where (21) places it.

The coherence-band reading. The numerical value (21) is *suggestive*, not derived. VERSF does not claim that (21) equals v = 246 GeV. The honest structural claim is weaker but more robust:

The electroweak scale lies within an emergent *coherence band* delimited by closure–CCC geometric-mean relations. Membership in this band does not require exponential fine tuning; it is the natural consequence of UV–IR pairing between two substrate-derived scales.

A coherence band of width one to two orders of magnitude around (21) easily accommodates v. The remaining task is not to derive 246 GeV from the geometric mean alone — which would require matching prefactors absent from (21) — but to identify the additional substrate parameter (likely a TPB-rate ratio) that selects v within the band. What can be said now is that the band edge sits near 5–6 TeV, and that no new physics has been observed up to current LHC sensitivities (~ 2 TeV for techni-resonances; ~ similar for KK modes) — a null result that is consistent with a smoothly-suppressed coherence-band edge of the kind §4 predicts, and that disfavours the alternative scenarios (composite Higgs, low-scale SUSY) which predict resonance structure at or below the same scale.

The structural point is preserved: on the closure–CCC ladder, a separation of seventeen orders of magnitude between v and M_P is one geometric-mean step plus a band-width adjustment, not thirty-two orders of fine tuning. The ladder reframes the magnitude of the hierarchy as a logarithmic, not exponential, problem.

---

9. Why the Standard RG Argument Misfires

We can now restate cleanly why the conventional naturalness argument fails under the VERSF reading.

The Wilsonian flow

$dm_H^2/d \ln \mu = \beta_m(g, m_H, \ldots)$ (22)

is a statement about a fixed ontological sector — the record layer — and is well-defined while μ remains within that sector's domain of validity, i.e. while $\mu < k_C$. The naive integration to M_P assumes the integrand is meaningful across the full interval. Under VERSF, this assumption fails for three concurrent reasons:

1. By the OSP, the inputs to β_m are weighted by $F_{adm}(k)$ and vanish for $k > k_C$; equation (4) is therefore physically meaningless beyond k_C even though it remains a well-defined formal integral.
2. By the CNPP, no record-layer propagation exists near M_P; the integrand has no degrees of freedom to flow.
3. By the admissibility kernel, the loop integrals (5)–(7) are physically truncated at k_C through $F_{adm}(k)$, so even within $\mathcal{L}_R$ the integration domain never reaches M_P.

The conventional argument conflates a legitimate flow within $\mathcal{L}_R$ (up to k_C) with an illegitimate continuation *through* the closure threshold into $\mathcal{L}_C$. The quadratic sensitivity to M_P is generated entirely by the illegitimate continuation and is, additionally, regulator-dependent in interpretation (Section 5). Once the category distinction is enforced and a physically grounded regulator is used, the catastrophe disappears.

---

10. Why Naturalness Looked Compelling

The naturalness argument was not a mistake at the level at which it was developed. Within the record layer — which is the only layer to which the renormalisation-group framework directly applies — Wilsonian extrapolation has been a consistently successful guide. Effective field theory has correctly predicted decoupling, scale-dependent couplings, and the breakdown of low-energy descriptions at expected thresholds across condensed matter, nuclear, and particle physics. The argument that an unprotected scalar should be sensitive to high-scale physics is rigorous *within record-layer dynamics*. There is no quarrel here with naturalness as a research heuristic for situations in which both endpoints of the extrapolation lie within a single dynamical layer.

The hierarchy problem became dominant for two specific reasons. First, the Higgs is the only known elementary scalar, and this concentrated naturalness pressure on a single state rather than distributing it across a sector. In extended scalar models — two-Higgs-doublets, MSSM Higgs sectors, supersymmetric extensions with many scalar superpartners — quadratic sensitivities are shared across several fields, and the requirement that any particular one be light is correspondingly weaker. The Standard Model has no such cushion: the same Higgs that gives masses to all other particles is also the only scalar whose mass could in principle be technically natural. Whatever UV-sensitivity the framework attributes to elementary scalars therefore falls on this one state, with nothing to absorb or distribute it. This sharpens the naturalness pressure into a near-singular focus on the Higgs and makes the seventeen-order gap feel acute in a way it would not in a richer scalar sector.

Second, in the decades when grand-unified-theory and superstring approaches to UV completion were ascendant, it was natural to assume that the Planck scale was inhabited by record-layer-style degrees of freedom — strings, branes, Kaluza–Klein modes, GUT gauge bosons — to which the standard extrapolation would apply. Under those assumptions, the seventeen-order gap really would have implied profound fine tuning, and the supersymmetric and composite responses were entirely reasonable replies.

The VERSF claim is therefore targeted, not sweeping. It does not deny that naturalness governs flow *within* the record layer. It denies that the upper limit of (4) lies within the record layer at all. The argument breaks only at the closure boundary — not generally, not for ordinary effective-field-theory situations, and not for the many domains (asymptotic freedom of QCD, decoupling of heavy quarks, hierarchies of nuclear scales) where Wilsonian reasoning has earned its place. What is being rejected is the *assumption that the Planck scale is a record-layer scale*. The naturalness puzzle dissolves precisely because the conventional argument's premise about the Planck scale is the one VERSF identifies as categorically incorrect; the rest of the EFT framework remains intact and operative.

Two consequences follow. First, VERSF makes contact with the bulk of successful EFT practice without disturbance: the framework neither supersedes nor displaces effective field theory within its proper domain. Second, the failure mode VERSF identifies is structurally specific. It applies to extrapolations that cross categorical layers; it does not apply to extrapolations within a single layer. This is why naturalness has not failed in the many domains where it works, and why it nonetheless misfires for the Higgs in the specific case that has dominated the hierarchy literature.

---

11. Distinguishing Predictions

If the hierarchy problem is a category error rather than a missing symmetry, the experimental landscape changes in identifiable ways.

Predictions are marked **[sharp]** if they admit direct experimental falsification in identified channels, or **[structural]** if they characterise the framework's posture rather than predict a measurable outcome. Sharp predictions are listed first.

11.1 Smooth turn-off rather than resonance tower. [sharp] Composite-Higgs models predict a tower of techni-resonances at the compositeness scale $\Lambda\_TC \sim 10$ TeV: identifiable, narrow, gauge-charged states accessible to collider searches. VERSF predicts *no* such tower at k_C. Instead, scalar-sector form factors should display a smooth suppression governed by F_adm(k), with the sharpness of the turn-off set by the index n in (9′). The discriminator is qualitative — presence vs. absence of resonance structure — and is already operative at the LHC.

This now connects directly to the §8 geometric-mean reading. With E_geo ≈ 5–6 TeV identified as the coherence-band edge under option (iii) of §4.3, the prediction is that no resonance tower should appear in the few-TeV region: instead, scalar-sector observables should show smooth suppression as one approaches the band edge from below. The current LHC null results — no techni-resonances up to ~ 2 TeV, no KK modes up to similar scales, no light SUSY partners — are exactly what this picture expects. Composite Higgs and low-scale SUSY accommodate the same null results at the cost of pushing their natural scale upward, weakening the naturalness motivation that produced them in the first place; under VERSF the absence of a tower at these scales is simply what the framework predicts. The framework is silent on whether SUSY partners, KK modes, or other beyond-SM structure exists for reasons independent of naturalness

— it merely does not require any such structure as a naturalness remedy, and so does not take credit for the broader pattern of BSM null results beyond its own predictive scope.

11.2 Geometric-mean ladder. [structural / sharp] If UV–IR coupling generates a ladder of physical scales (§8), the EW scale should not be the only intermediate rung. The *existence* of additional rungs along the closure–CCC ladder is a structural prediction; their *specific locations* depend on substrate parameters not yet derived from the master action. Conjectured axion-like-particle scales and as-yet-unobserved condensate phenomena are candidate rungs, but the framework does not currently specify which rung they should occupy or how the prediction differs from existing scenarios (e.g. seesaw mechanisms for neutrino masses, which already invoke high-scale physics). A sharp instance would require either deriving the next rung's location from the substrate or identifying a phenomenon that occurs at a geometric-mean rung and at no other naturally-motivated scale. The trichotomy in §4.3 — sector-specific ξ, multiple closure scales, or geometric-mean rung — compounds this under-determination: until the master-action derivation fixes which ξ governs which sector, ladder-rung locations are correspondingly free, and the prediction's sharpness is conditional on closing that derivation. Until then, the prediction is best read as: the EW scale is one of a class of intermediate coherence scales, not a singular point.

11.3 Correlations among mass scales. [sharp] Particle masses, in this picture, are eigenvalues of $\mathcal{C}[\rho]$ expanded around symmetry-broken vacua. Mass *ratios* should therefore satisfy non-trivial substrate-determined relations rather than being independent free parameters. Identifying such relations is a concrete programme of comparison with the SM mass spectrum and CKM/PMNS mixing data — partially advanced in the VERSF flavour papers.

11.4 Substructure in Higgs couplings. [sharp, weakly] If the Higgs is a coherence eigenmode of the record field, its couplings should display deviations from elementary-scalar predictions at energies approaching the coherence-stability bound near k_C. Precision Higgs programmes (HL-LHC, FCC-ee, CEPC, ILC) targeting per-cent-level coupling measurements could reveal these as small, *correlated* departures across multiple production and decay channels. The prediction is sharp in principle but weak in current sensitivity: the deviation pattern is not specified quantitatively without the master-action derivation of $\mathcal{C}[\rho]$.

11.5 No required low-scale UV completion. [structural] There is no structural reason to expect supersymmetric partners, composite resonances, or extra-dimensional Kaluza–Klein states at or near the electroweak scale. Their absence at the LHC is consistent with — not anomalous under — the VERSF reading. This is a non-prediction in the strict sense: it predicts the absence of features that other frameworks predict the presence of.

11.6 Decoupling of gravity from scalar mass. [structural] Because the Planck scale is a closure threshold rather than a UV sector of the record layer, gravitational physics does not directly renormalise the scalar mass. This is consistent with current bounds on Higgs–graviton couplings, and follows almost definitionally from the layered ontology.

---

12. Epistemic Stratification

Honesty requires separating what is established within VERSF, what is conditional, and what remains conjectural.

Proved (within VERSF).

- The two-layer ontology (closure vs. record) follows from the VERSF primitives of finite distinguishability and irreversible commitment.
- The OSP follows from the construction of $\mathscr{R}$_RG from admissibility-weighted loop integrals: its inputs vanish above k_C by the kernel property (8).
- The CNPP follows from the propagation conditions (C1)–(C3) together with their alignment with the standard QFT requirements of pole structure, unitarity, and locality.
- Admissibility closure regulates record-layer correlation functions via a suppression kernel of the form (7)–(8).
- The κ-field memory kernel K(τ) and its asymptotic form K_eff(τ) ~ A cos(ω_κ τ + φ)/τ are derived from the κ-field Lagrangian without phenomenological input (Taylor, *The Memory Kernel from First Principles*).
- The κ-field mass m^2_κ = λ_eff ξ^{-2} = (4/3) ξ^{-2} is a theorem of the K=7 minimal fact architecture and the CCC threshold (Taylor, *The κ-Field Mass as the Physical Hessian of the Commitment Constraint Surface*).
- BCB and TPB constraints are incompatible with naive continuum RG flow through the closure threshold.

Conditional.

- That the identification k_C ≡ m_κ is physically correct. The κ-field is the natural universal mediator of record-layer coherence, which motivates the identification, but no VERSF paper currently proves it as a theorem at the level of the master action. Pending such a proof, k_C ≡ m_κ should be read as a structural identification rather than a derived consequence of the upstream κ-field-mass result. Under option (iii) of §4.3, even this structural identification is replaced for the EW sector by k_C ≡ E_geo ≈ 5–6 TeV, with m_κ acting as the IR input to E_geo rather than as the regulator threshold itself.
- That the relevant ξ for the electroweak hierarchy is the one for which √(4/3) · ℓ_P / ξ ~ k_C / M_P at the EW scale. The cosmological CCC scale ξ ≈ 8.2 × 10^{-5} m gives m_κ ≈ 2.8 meV (fourteen orders of magnitude below the EW scale), so either a sector-specific ξ governs the EW hierarchy or the EW scale occupies a non-trivial rung on the geometric-mean ladder of §8 (where E_geo ≈ 5–6 TeV places v within one order of magnitude of the predicted band).
- That the coherence functional (17) yields a vacuum at ⟨ρ⟩ corresponding to v ≈ 246 GeV with the correct gauge-charge structure.

Conjectural.

- That the Higgs boson is, specifically, the lowest-lying coherence eigenmode of the record field rather than a fundamental scalar.

- That the geometric-mean relation (20)–(21) selects v within an order of magnitude with coefficients derivable from substrate parameters.
- That a full ladder of intermediate scales populates the closure–CCC interval and corresponds to observable particle masses.

The category-error thesis itself — that conflating closure and coherence layers is the underlying source of the naturalness puzzle — is robust across these strata. Even if particular numerical claims fail, the structural reframing stands or falls on whether the two layers are genuinely distinct dynamical categories. VERSF's pregeometric construction provides positive grounds for thinking they are.

---

13. Relation to Other Approaches

The VERSF reading shares structural features with several existing strands of work without reducing to any of them.

- **Asymptotic safety** seeks a non-trivial UV fixed point that renders gravity predictive at all scales. VERSF agrees that the standard Wilsonian extrapolation breaks down near the Planck scale but offers a different reason: not a fixed point in a single sector but a transition between sectors.
- **Emergent gravity** (Verlinde, Padmanabhan, Jacobson) treats gravity as a thermodynamic or entropic phenomenon. VERSF agrees that gravitational structure emerges from substrate-level commitment statistics; the present paper extends the emergence reading to the Higgs sector.
- **Composite Higgs** is sharply distinguished in Section 7: it is a within-field-theory emergence, whereas VERSF is an emergence from a pre-field closure substrate.
- **Relaxion mechanisms** dynamically select a small Higgs mass from cosmological evolution. These are complementary: a substrate that admits coherence locking might naturally implement relaxion-like dynamics as a record-layer effective theory.
- **Dimensional-regularisation naturalness** (Bardeen and others) observes that quadratic divergences are scheme-dependent and that the only physical sensitivity is logarithmic. VERSF endorses this observation (Section 5) and supplies the physical regulator that makes it concrete.
- **Cross-sector corroboration.** The K=7 architecture invoked here to fix k_C is not single-purpose machinery. Most strikingly, the same architecture supplies the CMB scalar spectral index $n_s \approx 1 - 2/N_\star \approx 0.964$ at $N_\star = 55$ e-folds via logarithmic κ-field displacement growth during inflation (Taylor, *The Spectral Density of the Commitment-Event Bath in the VERSF Framework*), comparable to the Planck 2018 value $n_s = 0.965 \pm 0.004$ — a percent-level agreement on an independently measured cosmological observable derived from the same K=7 architecture being deployed in this paper for the hierarchy. The §8 geometric-mean reading developed here is itself a second contact-with-data result from the same architecture: $E_{geo} \approx$ 5–6 TeV from K=7 + CCC + the observed cosmological vacuum-energy density, within one order of magnitude of v = 246 GeV, with the smooth-turn-off phenomenology compatible with current LHC null results.

> Two empirical contacts (a CMB spectral index and a TeV-scale coherence-band edge) issuing from one substrate architecture is meaningfully stronger evidence than either result alone. The same architecture also supplies the non-Markovian decay kernel governing radioactive nuclei (Taylor, *κ-Field Wave Dynamics, Geometric Memory, and Non-Markovian Decay*) and the gravitational sourcing of spacetime curvature through commitment-density gradients (Taylor, *Gravity from Tensorial Closure of Record Dynamics*). The hierarchy-problem reading developed here rests on a constraint package with independent applications, not on apparatus introduced solely to address naturalness.

The VERSF programme is not a competitor to a single proposal but a re-classification of what kind of problem the hierarchy is.

---

14. Conclusion

The hierarchy problem, in its standard form, presupposes three things: that the Planck and electroweak scales are commensurable, that the renormalisation group can flow between them, and that the Higgs is a fundamental scalar. All three presuppositions fail under VERSF.

The Planck scale is a closure threshold of the substrate, not a UV completion of the record layer. The renormalisation-group operator is constructed from loop integrals weighted by the admissibility kernel; its naive extension to the closure threshold is physically meaningless beyond k_C, even where the formal integral remains well-defined (Ontological Separation Principle). Closure-layer structure cannot propagate within the record layer because the propagation conditions (C1)–(C3) — and their QFT counterparts of pole structure, unitarity, and locality — fail at the threshold by construction (Closure Non-Propagation Principle). Admissibility closure regulates record-layer correlation functions via a physical suppression kernel that terminates loop integrals at a closure scale k_C far below M_P — the scale at which Causal–Coherence Compatibility fails for sustained record transport, set heuristically by the bit-budget cost of coherence relative to substrate ticking. And the quadratic divergence on which the standard argument rests is, in its physical interpretation, regulator-dependent rather than scheme-independent: it vanishes in dimensional regularisation and lacks unambiguous physical content there. The VERSF position is stronger than the dimensional-regularisation reply alone, because it supplies a physical regulator from substrate primitives rather than arguing that one scheme is more physical than another.

The Higgs, on this reading, is not a fundamental scalar exposed to Planck-scale fluctuations. It is a coherence eigenmode of the record field — built on the commitment density and shaped by a substrate-determined functional $\mathcal{C}[\rho]$ that includes record-current transport and admissibility coupling as constitutive terms. Its mass is not a bare parameter but a derived eigenvalue.

The naturalness problem arises only if one assumes that all scales belong to a single dynamical ontology. VERSF denies this assumption. The hierarchy is then not a failure of nature to protect the Higgs, but evidence that the Planck scale is not part of the Higgs sector at all.

What remains is the constructive task: deriving the explicit form of $F_{adm}(k)$ and $\mathcal{C}[\rho]$ from the VERSF master action, demonstrating that the resulting vacuum reproduces SM phenomenology in detail, and identifying further rungs of the closure–CCC ladder. That task is substantial, but it is *constructive* rather than restorative. There is no broken naturalness to repair.

Seventeen orders of magnitude between v and M_P, on the VERSF reading, are not a crisis. They are the structural signature of a layered reality whose layers were never the same kind of thing to begin with.

References

VERSF programme papers

- Taylor, K. *Admissibility Closure in VERSF.* [Establishes that record-layer correlation functions are regulated by a substrate-derived suppression kernel; the source of the kernel property (8) used throughout §4.]
- Taylor, K. *Causal–Coherence Compatibility and the Fact-Production Threshold.* [Defines $\xi = (\hbar c/\rho)^{1/4}$ and the CCC condition.]
- Taylor, K. *Causal–Coherence Compatibility and the Emergence of the Mesoscopic Scale.*
- Taylor, K. *A No-Go Theorem for Non-Simplicial Relational Substrates – K=7.* [Establishes K=7 minimal fact architecture via the Fano plane PG(2,2).]
- Taylor, K. *Derivation of the κ-Field Mass from Minimal Fact Architecture.* [Derives $\lambda_{eff} = 4/3$.]
- Taylor, K. *The κ-Field Mass as the Physical Hessian of the Commitment Constraint Surface.* [Proves $m^2_\kappa = \lambda_{eff} \cdot \xi^{-2}$ as a theorem, with Uniqueness of Constraint Penalty (§4), No-Alternative-Scaling Lemma (§5), and Constraint-to-Mass Proposition (§6).]
- Taylor, K. *The Memory Kernel from First Principles: A Dynamic Derivation of K(τ) in VERSF.* [Derives the κ-field memory kernel $K(\tau) \sim A \cos(\omega_\kappa \tau + \varphi)/\tau$ from the κ-field Lagrangian.]
- Taylor, K. *On the Equivalence of the κ-Field and the Commitment Field: A Unique Klein–Gordon Structure from Irreversible Fact Formation in the VERSF Framework.* [Lagrangian Uniqueness Theorem; field identification $s \equiv \kappa$.]
- Taylor, K. *The Origin of the Closure Scale in the VERSF Framework.*
- Taylor, K. *Coherence Scale, Memory Kernel Reduction, and Non-Markovian Gravity in VERSF.*
- Taylor, K. *κ-Field Wave Dynamics, Geometric Memory, and Non-Markovian Decay in VERSF.* [Derives the source-geometry dependence of the κ-field response: 3D-ensemble kernel $\sin(m_\kappa \tau)/m_\kappa$ exactly, with Lorentzian Fourier transform; worldline-tube kernel $\cos(m_\kappa \tau)/\tau$ asymptotically. Relevant to the §4.1 grounding of $F_{adm}(k)$.]
- Taylor, K. *The Spectral Density of the Commitment-Event Bath in the VERSF Framework.* [Derives the κ-field bath spectral density from K=7 mode structure; establishes white-noise universality class ($\alpha = 1$) in the IR; yields CMB scalar spectral index $n_s \approx 1 - 2/N_\star \approx 0.964$ for $N_\star = 55$ e-folds.]

- Taylor, K. *Fold, Fact, Planck, and Closure: Structural, Realization, Metric, and Coherence Thresholds in VERSF.* [Distinguishes the four substrate-level thresholds — structural (fold), realisation (fact), metric (Planck), and coherence (closure) — used in §§2–3 to motivate the two-layer ontology.]
- Taylor, K. *The Two-Planck Framework.*
- Taylor, K. *A Hidden "Middle Scale" in the Universe.* [Mesoscopic CCC coherence scale $\xi \approx 8.2 \times 10^{-5}$ m.]
- Taylor, K. *Gravity from Tensorial Closure of Record Dynamics.*
- Taylor, K. *The Constitutive–Predictive Bridge in VERSF.*
- Taylor, K. *K=7 Convergence: Six Independent Routes to the Constraint Count.*
- Taylor, K. *Derivation of the Standard Model Gauge Group within VERSF.*
- Taylor, K. *CKM/PMNS Flavour Mixing in VERSF.*
- Taylor, K. *The κ-Field Origin of* $|\psi|^2$*: Rate Law, Quadratic Measure, and Uniqueness.*
- Taylor, K. *Admissibility Closure and the Uniqueness of Physical Entropy: A Derivation of* $\eta = 1$ *in the VERSF Framework.* [Parallel application of the admissibility methodology used in this paper: derives the entropy-partition identification $\eta = 1$ and commitment barrier $\Phi_c = 1$ as theorems via a No-Alternative Partition Theorem. Cited as evidence of the methodology's reach across the programme rather than as an input to the hierarchy argument.]

Selected external references

For naturalness, the hierarchy problem, and dimensional-regularisation perspectives:

- 't Hooft, G. (1980). *Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking.* In *Recent Developments in Gauge Theories*, NATO Advanced Study Institutes Series. [The canonical formulation of technical naturalness.]
- Bardeen, W. A. (1995). *On naturalness in the Standard Model.* Fermilab preprint FERMILAB-CONF-95-391-T. [The dimensional-regularisation argument that quadratic divergences lack scheme-independent physical content.]
- Giudice, G. F. (2008). *Naturally speaking: the naturalness criterion and physics at the LHC.* In *Perspectives on LHC Physics*, World Scientific. [Pre-LHC framing of the naturalness expectation.]
- Dine, M. (2015). *Naturalness under stress.* Annual Review of Nuclear and Particle Science 65, 43–62. [Post-LHC reassessment.]

For composite-Higgs models (the closest field-theoretic relative of the emergence reading):

- Contino, R. (2011). *The Higgs as a composite Nambu-Goldstone boson.* In *Physics of the Large and Small (TASI 2009)*, World Scientific. [Standard pedagogical review.]
- Panico, G. and Wulzer, A. (2016). *The Composite Nambu–Goldstone Higgs.* Lecture Notes in Physics 913, Springer. [Comprehensive treatment.]

For prior observations of the $\sqrt{M_P \cdot M_\Lambda} \sim$ TeV geometric-mean coincidence (cited in §8):

- Dirac, P. A. M. (1937). "The cosmological constants." *Nature* 139, 323. [Large-numbers hypothesis proposing structural relations between cosmological and microphysical scales.]
- Dirac, P. A. M. (1938). "A new basis for cosmology." *Proc. Roy. Soc. A* 165, 199–208.
- Cohen, A. G., Kaplan, D. B., and Nelson, A. E. (1999). "Effective field theory, black holes, and the cosmological constant." *Phys. Rev. Lett.* 82, 4971–4974. [arXiv:hep-th/9803132. Derives a TeV-scale geometric-mean cutoff from black-hole-entropy bounds on effective field theories; cited in §8 as a prior observation of the $\sqrt{(M_P \cdot M_\Lambda)} \sim$ TeV coincidence.]
- Banks, T. (2000). "Cosmological breaking of supersymmetry?" *Int. J. Mod. Phys. A* 16, 910–921. [arXiv:hep-th/0007146. Discusses cosmological-constant scaling within holographic and SUSY-breaking frameworks.]

For emergent-gravity and entropic perspectives:

- Jacobson, T. (1995). *Thermodynamics of spacetime: the Einstein equation of state.* Physical Review Letters 75, 1260.
- Padmanabhan, T. (2010). *Thermodynamical aspects of gravity: new insights.* Reports on Progress in Physics 73, 046901.
- Verlinde, E. (2011). *On the origin of gravity and the laws of Newton.* Journal of High Energy Physics 04, 029.

Full VERSF corpus: versf-eos.com