The Vacuum as Weather

One Machine, Two Generators, and the Single Knob at the Bottom

J. Konstapel · Leiden
Third edition — 28 July 2026

I. A machine with no designer

Roughly every four years the equatorial eastern Pacific warms by a few degrees. Peruvian fishermen named it after the Christ Child because it peaked near Christmas. It is now the most closely instrumented natural oscillation on Earth, and its internal workings are known well enough to be written down.

Here is what happens. A warm anomaly in the eastern Pacific weakens the trade winds. The weakened winds change the slope of the thermocline, deepening it in the east, which warms the surface further — a positive feedback that amplifies the anomaly that caused it. But as the anomaly strengthens, the ocean transports heat away from the equator. The thermocline shoals; the warm anomaly is undercut and dies. The system swings into its opposite phase, cold water in the east, easterly wind anomaly, convergent heat transport back toward the equator — and the heat content recharges. Then it happens again.

This is Jin's *recharge oscillator*, and it has a property worth pausing on. Nothing in it is regulating anything. There is no controller, no set point, no external clock beating four years. The period emerges from the transit times of oceanic waves and the rate at which heat accumulates and discharges. The regularity is a consequence of the coupling, not of a rule imposed on it. Aristotle's efficient cause closes inside the system; no external prime mover is required, and inserting one explains nothing.

That is the first thing to hold onto: the Pacific is a machine, and it is a machine without a designer.

The Zebiak–Cane model captures this in equations coupling a reduced-gravity ocean, a mixed layer with an advection–diffusion temperature equation, and a Gill–Matsuno atmosphere. It has a single dial, the coupling strength μ between ocean and atmosphere. And what that dial does is the second thing to hold onto.

Below a critical value μ_c, the system is damped. Perturb it and the oscillation decays. Above μ_c, at a Hopf bifurcation, the same system oscillates on its own, indefinitely, with a spectral peak near four years. Nothing else changed — not the physics, not the equations, not the constants. Only the coupling.

This is the concrete meaning of a claim that otherwise sounds merely provocative: *what counts as the law of a system is a function of its context*. The Pacific does not “have” the law of damped return to equilibrium, nor the law of self-sustained oscillation. It has whichever of those the local coupling puts it in reach of. The law is a fixed point of the dynamics under the prevailing conditions, not a property engraved beneath them.

And a third thing, subtler and more important than it looks. In the subcritical regime, below the threshold, adding wind-stress noise produces an ENSO-like signal anyway. The oscillation is latent in the relaxed state; noise does not create it, it excites what was already available. Order does not have to be built up out of disorder. It is closer to the truth to say that the free, unconstrained state contains the possibility, and that structure is the tensed exception drawn out of it.

II. Grammar versus identity

The Pacific is not the only system that works this way. The same equations — a positive feedback, a delay, a threshold, a storage-and-release cycle — describe a laser above threshold, a beating heart, a neuron, a chemical clock, a stick-slip fault, a triode oscillator. Balthasar van der Pol built the circuits at Philips in Eindhoven in the 1920s and heard, in a telephone headset, the irregular hiss just before the frequency jumped: the first human observation of deterministic chaos, made with an earpiece.

Why do such unlike things obey the same equations? Because near a bifurcation almost all of the substrate washes out. What survives is a *normal form* — the minimal equation compatible with the symmetries of the transition. This is proven mathematics, not analogy.

But it cuts both ways, and the cut is the reason most cross-scale claims fail. If the vacuum and the Pacific Ocean share a normal form, that resemblance by itself tells you almost nothing about the vacuum. It tells you only that both sit near a Hopf threshold. The very universality that produces the similarity is what drains the similarity of information.

So there are two distinct claims, and they must not be run together:

The grammar claim. Tension–release machines with thresholds are a generic class, and both climate and vacuum belong to it. This is close to established; the ENSO literature supplies one worked example and the Sakharov–Volovik tradition supplies the vacuum-as-elastic-medium reading.

The identity claim. These are not merely similar machines but *the same machine at different layers of one nested system*, such that the layers are related by a definite recursion.

The identity claim is the interesting one. It is also the one that has historically dissolved into unfalsifiable resemblance-hunting. The remainder of this essay is an attempt to state it so that it does not.

One further distinction earns its place here. Ordinary scale invariance — a power law — says the system looks the same under *any* magnification. That is a weak and nearly unfalsifiable claim. *Discrete* scale invariance says the system reproduces itself only under specific magnifications, at ratios λ, λ^2, λ^3. This is a much stronger statement, and it leaves a measurable signature: log-periodic corrections riding on the power law, described by complex critical exponents. Sornette and collaborators found these in rupture acoustics, earthquake precursors, two-dimensional turbulence, and financial bubbles; Zhou and Sornette detected more than twenty harmonics of a single logarithmic frequency in high-Reynolds-number turbulence, corresponding to a preferred scaling ratio close to 2.

If layers exist, in other words, there is a number λ that says how far apart they are, and it is measurable.

III. Inverse renormalization: stripping down to the zero layer

The standard renormalization group runs one way: start with microscopic detail, coarse-grain, and watch the irrelevant couplings die until only a few effective parameters remain. Physics has spent fifty years walking down that road, and it has worked.

Run it backwards and something else appears. Start with a fully organized system — an ecology, a climate, a star — and strip organization rather than resolution. Remove the biota; the atmosphere remains, still cycling. Remove the atmosphere; the magnetosphere remains, still storing and discharging. Remove the planet; the solar wind remains, a turbulent magnetized flow with a well-measured inertial range. Remove the plasma; the electromagnetic field remains. Remove the field's excitations, and what is left is not nothing — it is the medium that still has a stiffness, still supports waves, still has a characteristic impedance.

That residue is the vacuum. Note what has been stripped: *organization*, not matter. The vacuum on this reading is not the empty box in which layers sit. It is the limit of the stripping process — the zero layer, the state that remains when every level of structure has been removed and there is nothing left to remove.

This is not a new intuition. Sakharov's induced-gravity proposal treated curvature as the elastic response of the vacuum, gravity emerging as the medium's resistance to being deformed rather than as a fundamental interaction. Volovik developed a sustained analogy between superfluid helium-3 and the quantum vacuum, in which the "fundamental" excitations of the low-energy theory — chiral fermions, gauge fields, an effective metric — appear as collective modes of a medium whose real structure lies above the effective theory's cutoff and is invisible to observers made of those modes.

That last clause is the hinge of everything that follows. If observers are made of the medium's excitations, there are things about the medium they cannot see from inside.

IV. The single knob

How much structure can an internal observer actually read off the vacuum?

The linear electromagnetic response of free space is fixed by two constants: the permittivity ε_0 and the permeability μ_0. Every other electromagnetic property of the vacuum is a combination of these two. There are exactly two independent combinations.

The product gives a wave speed:

$c = 1 / \sqrt{(\varepsilon_0 \mu_0)}$

The ratio gives a wave impedance:

$Z_0 = \sqrt{(\mu_0 / \varepsilon_0)} \approx 376.730\ \Omega$

Speed and impedance: in mechanical terms, this is a medium characterized by a stiffness and an inertia — and a medium with two independent moduli has internal structure, because the two can vary separately and their variation is a fact about what the medium is made of.

Except that the vacuum does not, in the end, have two.

Since the 2019 revision of the SI, *c* is fixed by definition. It is a unit conversion between the metre and the second, not a measurement outcome. Fixing it does not discover anything about the vacuum; it decides what "metre" shall mean. So of the two combinations, one has been legislated into a convention. That leaves the impedance.

And the impedance is not free either. It is locked to the fine-structure constant:

$\alpha = e^2 Z_0 / (2h)$

Substituting the measured values gives $\alpha = 7.2973525 \times 10^{-3} = 1/137.035999$. Exactly — this is an identity, not a numerical coincidence. In the post-2019 SI, where *e* and *h* are also defined constants, the relation runs the other way and it is μ_0 that inherits an experimental uncertainty, entirely inherited from α.

Follow this through. **The complete measurable structure of the vacuum — everything an observer made of light and charge can read off it from inside — compresses to one dimensionless number.** Not a stiffness and an inertia independently. Not two moduli. One knob.

Three consequences follow, and they are the substance of this essay.

First, this explains why the vacuum looks empty. A medium whose internal observers can measure only a ratio, never the density and the stiffness separately, is indistinguishable from nothing — except through the value of that one ratio. The structure is not absent; it is *underdetermined for observers embedded in it*. That is not a philosophical remark but a consequence of dimensional analysis: the units in which a stiffness and an inertia would be separately expressed are themselves constructed out of the medium's own behaviour.

Second, it prunes the field of candidate models drastically. Any proposal requiring two independent vacuum properties is, from inside, empirically indistinguishable from a proposal requiring one. The vacuum carries a single degree of freedom. Anyone offering a structure for it therefore owes exactly one thing: the mechanism that fixes 1/137.036. Everything else is decoration on an unobservable.

This is a real narrowing. It disqualifies nothing on grounds of implausibility and everything on grounds of redundancy — which is a cleaner blade. It also happens to favour the sparest possible ontology: one self-resonating entity rather than a substance with parts. Wheeler's one-electron universe was reached by a different route and abandoned for the wrong reason (it was framed in particle language at a time when de Broglie's wave reading had not yet been absorbed), but the topology it gestures at survives the pruning better than the alternatives.

Third, it closes a circle with the topological reading of charge. Williamson and van der Mark proposed that the electron is a photon confined to a toroidal path — charge as a topological winding number of light circulating in the medium, mass as trapped circulation rather than as substance. On that reading $\alpha = e^2 Z_0/2h$ says something specific: *the impedance of the medium and the winding number of its knots are not independent quantities*. One relation binds the sea to its vortices. Van der Mark later pushed the same logic on energy-balance grounds, tracking the fraction of internal kinetic energy from the Earth–Sun system ($\sim 10^{-14}$) through hydrogen ($\sim 10^{-8}$) and the deuteron ($\sim 10^{-2}$) to the proton (~ 0.37), and arguing that granularity must terminate where that fraction approaches unity — making the proton the smallest stable knot and the medium continuous below it.

Whether or not that specific hierarchy holds, the structural point stands: if the vacuum is a medium and particles are its stable circulation patterns, then the medium's one adjustable property and the particles' topology are the same fact seen twice.

There is an operational corollary that matters more than the interpretive ones. **Any local structure, tension, or context in the vacuum can only manifest as a local variation of α.** There is no second channel. This is why a measurement of $\Delta\alpha/\alpha$ in a strong gravitational potential — a white dwarf photosphere, say — is not one test among many. It is the only window there is.

V. The recursion lives in the ratios

If there are layers, something has to say how one layer generates the next. The most economical answer on offer comes from a direction that physics has not looked in, and it was reached from data rather than from theory.

Ray Tomes (ray.tomes.biz; Cycles Research Institute, Auckland) came to it by an unusual road. Trying to predict economic cycles by computer, he found that the first four periods he could isolate in the New Zealand economy were near-exact fractions of 35.6 years — one quarter, one fifth, one sixth, one eighth. Then, in forty-four years of weekly corn prices, he found a set of cycles whose frequencies stood in exactly the ratios of the white notes of the piano plus two black ones, with the original four forming a major chord. He did not mention it to anyone, on the reasonable grounds that it sounded absurd. Some years later he visited the Foundation for the Study of Cycles in the United States and discovered that Edward Dewey had found the same periods decades earlier, in entirely different series, from different countries, over different intervals. The disappointment of losing priority was outweighed by what the independent confirmation implied.

Dewey's own conclusions, after collating several thousand cycle reports, were four: everything that can be made into a time series has cycles in it; the same periods recur across unrelated fields; cycles sharing a period are phase-synchronized rather than merely coincident in length; and the ratios between commonly reported periods are dominated by small integers, especially 2 and 3. That compilation is a substantial empirical object, and it has never been properly explained.

Tomes's explanation is a single axiom:

> ***The universe consists of a standing wave which develops harmonically related standing waves, and each of these does the same.***

Read that carefully, because it is not the fractal one expects. It is not a shape repeating at different magnifications. It is a *recursion in ratio space*: each wave spawns waves at integer ratios to itself, and each of those does the same, so that the structure of the universe lives in the network of ratios between coupled rhythms rather than in any repeated form.

Nor is the recursion postulated. It follows from nonlinearity. Maxwell's equations under general relativity are a nonlinear system — energy warps the metric, so light's propagation depends on the light already present. In any nonlinear system, waves develop harmonics: the wave shape changes over time, which is another way of saying that integer-ratio components appear. In three dimensions these harmonics are themselves standing waves, so they do the same thing. The recursion is what nonlinearity does when it is allowed to run for long enough.

The mechanism that gives the theory its teeth is the accounting of *strength*. Energy flows from each wave to its harmonics, mostly to the low ones. But some harmonics can be reached by many routes and others by only one — the twelfth can be built as 12, 6×2, 2×6, 4×3, 3×4, 2×2×3, 2×3×2, 3×2×2, while the eleventh and thirteenth can be built only one way each. Harmonics reachable along many paths accumulate energy from all of them. So the strong scales of the universe are the ones with many routes into them, and the resulting pattern is not smooth: multiples of 12 stand out, strong harmonics come in families related by factors of 2, and across any octave the strongest harmonics fall on the intervals of just intonation. Between 48 and 96 the leading harmonics are 48–60–72–96, which is a major chord.

This is why the ancient phrase about the music of the spheres keeps returning without ever becoming respectable. The musical intervals and the strong scales of a nonlinear recursive medium are the same object, because both are simply the ratios reachable in the largest number of ways. Pythagoras used only 2 and 3; Galilei's father established that the major third is really 5/4; the theory says how often each prime should enter, and the answer matches musical practice, including the awkward status of the ratio 7 that Euler complained about and that Indian and blues music use anyway.

Two features of this deserve emphasis for the argument being made here.

First, **the recursion is top-down**. Structure forms from the largest scale downward, not from the smallest upward. Whatever one concludes about the cosmological consequences Tomes draws, this is the correct orientation for a layered picture in which the vacuum is the *residue* of stripping organization rather than the *building block* from which organization is assembled.

Second, **the ratio-space fractal is the right kind of fractal for what the ENSO material actually shows.** The one strictly fractal object in the ENSO literature is not a spatial pattern; it is the devil's staircase, the Cantor-like set of locking ratios between the system's internal frequency and the seasonal cycle. Zoom in anywhere on the ratio axis and you find further locking windows between the windows, indefinitely. Arnold tongues in the coupling–ratio plane are the same structure seen from another side, and above a critical nonlinearity the tongues overlap and the system goes chaotic by the quasi-periodicity route.

Nor is this confined to climate. The Kirkwood gaps in the asteroid belt are Arnold tongues; the same circle map governs them. The mathematics linking layers is not analogical here — it is literally the same equation appearing at the planetary and the orbital layer.

And it supplies a candidate for what the recursion between winding levels might be. In circle maps, the ratios most resistant to mode-locking — the last to be captured as nonlinearity increases — are the most irrational ones, the metallic means. A system driven by recursive coupling should be found sitting on those, not because of any numerological property but because they are the states that survive.

Tomes generates the rational ratios from above; mode locking selects them from below. In both cases the measure of a ratio's importance is the same: how many routes lead to it. That is the count of ways a harmonic can be built, and it is also what sets the width of an Arnold tongue. Two traditions that never met are measuring one quantity.

That convergence suggests an obvious bridge between the two research traditions of this essay: if strong harmonics and knot winding numbers are the same kind of object, particle mass ratios should sit on Tomes's strong harmonics. That bridge was tested, and it is not there. The result is given in §IX, together with the null hypothesis that makes it a result rather than an impression.

VI. The bottom rung, derived

Everything to this point has been architecture. This section is the first place where a number comes out.

A medium in which all the constituents share one collective state — a condensate — has a length it generates from its own properties, and does not get from outside. Push an obstacle into it and the density does not recover instantly at the edge; it takes a definite distance to heal. That distance is the **healing length**, ξ.

It has a definite value because two costs are in competition. Bending the collective wave sharply costs kinetic energy, so the medium wants to spread any disturbance over a long distance.

Changing the density costs interaction energy, so it wants to return to normal as quickly as possible. The balance of the two fixes one length.

What that length does is more important than what it is. It splits behaviour in two. Disturbances longer than ξ: the medium responds as a whole, with sound waves and a linear dispersion. Disturbances shorter than ξ: the medium responds as separate constituents, with a quadratic dispersion. The bend between the two regimes sits at ξ. This is the Bogoliubov spectrum, and it is the reason a condensate is not scale-free: it has one preferred scale, and the preference is generated internally.

That is exactly what the layer picture was missing. The Scale-Invariant Vacuum programme of Maeder and Gueorguiev forbids a preferred scale — its whole content is invariance under rescaling, and homogeneity forces its conformal factor λ to depend on time alone. Tomes asserts a preferred scale of $10^{4.5}$ but supplies no mechanism for its value. A condensate *produces* one.

Now put the vacuum's own quantities into it. The wave speed of the medium is c. The mass that belongs in it, following van der Mark's energy-balance argument, is the proton's — the point where granularity terminates. Then

$$\xi = \hbar / (\sqrt{2} \cdot m_p \cdot c) = 0.1487 \text{ fm}$$

against a measured proton charge radius of 0.84075 fm. A factor 5.65 apart.

That factor is not a discrepancy to be fitted away. It is an artefact of the $\sqrt{2}$ in the conventional definition of ξ. Remove it and what remains is startlingly clean:

$$\mathbf{r_p / \bar{\lambda}_p = 0.84075 / 0.21031 = 3.998}$$

Four, to six parts in ten thousand, where $\bar{\lambda}_p$ is the reduced Compton wavelength of the proton.

The status of that four must be stated precisely, because it is the difference between a derivation and a coincidence. **At present it is an observation, not a result.** To become a result it requires a topological stability condition: a demonstration that a closed three-dimensional electromagnetic field is stable at a winding of exactly four and not at two or six. Until that calculation exists, the bottom rung is anchored on a measured ratio rather than a derived one.

What can be derived, now, without any assumption, is the rung above it. From the proton's Compton wavelength to the Bohr radius:

- $\bar{\lambda}_e / \bar{\lambda}_p = 1836.2$ — the proton-to-electron mass ratio
- $a_0 / \bar{\lambda}_e = 137.036$ — exactly $1/\alpha$, since the Bohr radius *is* the electron Compton wavelength divided by α
- together: $\mathbf{a_0 / \bar{\lambda}_p = (m_p/m_e) \times (1/\alpha) = 2.516 \times 10^5 = 10^{5.401}}$

The step from the nucleon layer to the atomic layer is built out of precisely the two dimensionless numbers this essay has already privileged: the mass ratio of the two stable knots, and the coupling constant of the medium. Nothing was fitted and nothing was chosen.

It is also not Tomes's step. He gives 10^4.5; this rung is 10^5.40, a factor of eight away — nearly a full decade. That is a real disagreement, and it is the useful kind: either the nucleon-to-atom transition is one of his secondary scales rather than a major rung, or his major spacing does not hold at the bottom of the ladder. Both are checkable against his own published tables.

VII. Two generators, not one

The obvious next move is to iterate: multiply by 10^5.40 again and see what lands. That move fails, and the reason it fails is the most useful structural fact in this essay.

The generator of the first rung cannot produce the second. It worked because two ingredients were available: a mass ratio between two stable knots, and the coupling constant of the medium. One layer up, neither is available again. Chemistry runs on the same α and the same electron mass; no new coupling enters and no new stable mass ratio appears. And the molecular scale is correspondingly not 10^5.40 above the Bohr radius but a factor of one to ten above it — a molecule is a few atoms across.

The derivation stops at the atom, not because the arithmetic disagrees but because the ingredients run out.

That is not a failure of the layer picture. It locates a seam in it. Above the atom, scales are not set by ratios of constants at all; they are set by *balances* — thermal energy against surface tension, diffusion against reaction time, inertia against viscosity, self-gravity against quantum pressure. These are quantities of a different type. They arise from a competition, not from a division.

And the healing length is precisely a quantity of that type. It is the balance between bending cost and interaction cost. So the ladder plausibly has two generators:

Below the atom — discrete scales, fixed by ratios of absolute constants. Matter here is quantised; there are stable knots and nothing in between, so the steps are large and hard.

Above the atom — continuous scales, fixed by dynamic balances. Once chemistry exists, structure becomes fluid, and the characteristic lengths follow from competitions between forces rather than from constants.

The astrophysical evidence that the second generator is real and quantitative comes from an unexpected direction. In ultralight-boson dark matter models, the condensate's own internal length produces solitonic cores at the centres of galaxies on the order of a kiloparsec, and that core radius is measurable in rotation curves. A recent fit to the dwarf irregular DDO 168 within the Gross–Pitaevskii–Poisson formalism gives a core radius of 2.4 kiloparsecs. Whether or not ultralight bosons exist, the machinery demonstrates the principle: a medium whose own parameters fix one length, and that length visible as macroscopic structure. The Gross–Pitaevskii equations are indifferent to what the constituent is; they describe the fluid dynamics of a condensate and nothing more.

That the measurement method splits along the same seam as the mechanism is worth noticing. Below the atom there is no continuum to analyse — matter is discrete, and the only available approach is deduction from constants. Above the atom there are continuous, measured distributions across many decades, and the only available approach is spectral. Two halves, two generators, two methods. A seam that shows up independently in the physics and in the instrumentation is more likely to be real than one invented to rescue a table.

VIII. The ladder

With that in hand the ladder can be written down, but it must now be read as two documents stapled together rather than one.

Layer	The tension–release cycle running there
Vacuum (zero layer)	vacuum fluctuations; a single readable degree of freedom, α
Nucleon	self-closed circulation of light; scale $ƛ_p$, radius $4ƛ_p$
Atom	excitation and decay; scale $a_0 = ƛ_p \times (m_p/m_e) \times (1/\alpha)$
— seam: generator changes —	
Cell / dust scale	metabolic cycles, action potentials; also a common interstellar grain size
Moon / planet	**climate: ENSO, the recharge oscillator**
Star	the solar cycle (~11/22 yr), grand minima
Galaxy	density waves, star-formation cycles
Hubble scale	supercluster spacing

The two rungs below the seam are computed. Everything above it is observed, and the spacing there remains an open question rather than a derived quantity.

For completeness: iterating the derived step 10^5.40 upward from the proton's Compton wavelength produces 13 μm, 3.4 m, 840 km, 1.4 AU, 5.6 light years, 1.4 million light years. Those land within factors of 1.3, 2.0, 7.6, 1.4, 5.6 and 14 of a cell, a human, an Earth radius, an orbital radius, an interstellar separation and a galactic scale respectively.

This sequence should be treated with active suspicion, and the reason is instructive. The reference list against which the "hits" were scored was assembled by the author after computing the rungs. With steps of 5.4 decades and categories — cells, planetary orbits — that themselves span orders of magnitude, near-matches are close to guaranteed. The correct discipline is set out in §X. Until that test is run, the sequence is a pattern and not evidence, and §VII gives an independent physical reason to expect that it cannot be a recursion in any case.

IX. What we measured

Two layers have been tested, and the results are given here in full because the method matters as much as the numbers.

Data. The NINO3.4 sea-surface temperature anomaly series (NOAA PSL, HadISST1.1), monthly from January 1870 to April 2026, 1876 points; and the annual mean total sunspot number (WDC-SILSO) from 1700 to 2008, 309 points. Detrended fluctuation analysis, first order. Each series carries its own null: forty randomly shuffled surrogates, which must return an exponent of 0.50.

	characteristic cycle	inside the cycle	outside
Ocean (NINO3.4)	3 years	$\alpha = 1.43$ ($R^2 = 0.995$)	$\alpha = 0.58$ ($R^2 = 0.984$)
Sun (sunspot number)	11 years	$\alpha = 1.65$ ($R^2 = 0.996$)	$\alpha = 0.48$ ($R^2 = 0.902$)

Shuffled null: 0.49 ± 0.06 for the ocean, 0.49 ± 0.07 for the sun. The method is calibrated.

What the result says. Both layers show the same architecture, and it was not built in: a steep, strongly persistent branch inside the system's own cycle, a break exactly at the cycle duration, and beyond it an exponent indistinguishable from uncorrelated noise. Past its own cycle, neither system has memory. The machine charges, discharges, and forgets.

The outer exponents coincide — 0.58 and 0.48 against a noise of ±0.06. That is one number, not two. The inner exponents differ, 1.43 against 1.65, but the solar figure rests on six points at ten samples per cycle against the ocean's twelve per year; settling whether that difference is physical requires the monthly sunspot series.

What the result does not say. The two break points are not the fingerprint of one vacuum length. If they were, a single ξ would have to take two values, which it cannot — ξ is a property of the medium, not of what floats in it. Three years is the recharge time of the Pacific; eleven years is the solar dynamo. What the measurement establishes is that *any* medium with an internal scale carries this two-regime signature. That is the generic-form observation of §II again: the resemblance is real, and for that very reason it carries little information on its own.

And one thing that was tested and failed. The natural bridge between the two research lines of this essay — Tomes's strong harmonics and the knot models' winding numbers — would be that particle mass ratios sit on the strong harmonics. Twenty-one measured mass ratios, in units of the electron mass, were compared against the champion numbers of Tomes's ordered-factorisation count computed to 10^6. Median deviation from the nearest strong harmonic: 4.40%. For random numbers drawn log-uniformly over the same range: 4.55%. **$p = 0.44$.** Against Tomes's stricter criterion, the strongest harmonic per octave, the fit is worse than chance: 22.8% against 18.5%.

Individual near-misses look impressive — the kaon and the tau fall within 0.6% — and they are exactly what twenty-one random draws produce, because sixty-five champion numbers over six decades sit about a quarter apart in log space. That bridge does not exist.

It follows that the same discipline must be applied to Tomes's own structure-scale claim. He compares a predicted 10^4.56 with an observed 10^4.5 and reads it as agreement, without a null. Given the density of his strong harmonics, that null has to be computed before the agreement means anything. His tables are downloadable, so it can be.

X. What would settle it

Two questions remain open, one on each side of the seam, and each has a decidable form.

Below the atom: why four?

The factor $r_p = 4\lambda_p$ must either follow from a topological stability condition or be abandoned as a coincidence. This is a closed mathematical question requiring no new data, and the machinery for it already exists and is mature.

Faddeev and Niemi showed in 1997 that a nonlinear field theory admits stable knotted solitons — explicitly presented as the first real realisation of Kelvin's 1867 proposal that atoms are knotted vortices, an idea that had waited 130 years for a mechanism. In the Skyrme–Faddeev model these solitons are classified by an integer Hopf charge, which is a linking number. The spectrum has been computed: Battye and Sutcliffe constructed solutions up to charge eight and found the first genuine knot, the trefoil, at charge seven; Sutcliffe later reached charge sixteen and found a range of torus knots. At each charge there is a definite minimum-energy configuration, and the configurations change character as the charge rises — buckling at three, linked pairs at five and six, the trefoil at seven.

The question “why four and not two or six” is therefore not philosophical in this setting, and a first answer is already on the record. In the Skyrme–Faddeev model the minimum-energy configuration is a twisted torus — a single closed loop carrying winding — only up to Hopf charge four. From charge five upward the minimum-energy configurations are links, and from seven they are knots. **Four is the last charge at which the soliton is still one simple torus.**

If the proton is the largest stable single-torus configuration, four is therefore forced rather than chosen: two and three are also tori and are permitted, but six cannot be a simple torus at all. That is exactly the shape of selection rule the argument requires — an upper bound following from topology rather than from a fit.

Two cautions travel with it, and they set the remaining work.

The first is a category distinction that must not be blurred. The Hopf charge is a linking number; the 3.998 of §VI is a ratio of a radius to a wavelength. That both are four is striking, but they are not the same four until someone shows how one yields the other. That derivation is the missing step, and it is the whole of it.

The second is that the ground state at charge four is not settled: later minimisation runs report new ground-state configurations at charges four and five, and a symmetry-breaking deformation at six, with energies fitting the Vakulenko–Kapitanskii scaling E ~ |Q|^{3/4} more closely. Which shape wins at four is still in motion.

A further constraint narrows the search regardless. The Hopf charge is the product of two winding numbers, so a four is either 2×2 or 4×1; and in the magnetohydrodynamic analogue, torus-knot configurations exist only when the poloidal winding number is unity while the toroidal one may be any integer. If that selection rule carries over, the four is necessarily 4×1.

And there is one structural feature of this literature that closes the essay’s own theme rather than adding to it. The Skyrme–Faddeev model fixes the *shape* and the energy ratios of its solitons, but not their absolute size: that follows from the ratio of the two terms in the Lagrangian, which is an input. The knot literature therefore supplies the discrete topology and not the length. The condensate supplies the length and not the topology. The two generators of §VII reappear here on a single object, which is the clearest indication yet that the seam is real —

and it says where the four will have to come from: from the combination, not from either side alone.

Above the atom: is there a preferred spacing at all?

The test must avoid the selection bias that invalidated §VIII's sequence, and there is a construction that removes it rather than managing it. Weighting category widths does not help — every weighting is itself a free parameter, and the bias simply moves into the weights.

Instead, build a single continuous function: how much mass or structure exists per scale, plotted against log L. Then ask whether that function has periodic peaks in log L. This is a Fourier transform in logarithmic space, and it requires no boundary to be drawn and no example to be chosen. The null is built in: the power at the candidate frequency is compared with the power at every other frequency, and both are read off the same plot.

This is Sornette's log-periodicity test, and the argument has come full circle to it. Discrete scale invariance is detected exactly this way — a peak in the log-spectrum at the frequency corresponding to the preferred ratio.

The practical difficulty is that no single measured distribution spans femtometres to gigaparsecs. But several real, continuous, independently measured distributions each span many decades: the galaxy mass function, the halo mass function, the stellar initial mass function, the size distributions of asteroids and of interstellar dust. Stitched at scales where the measurements overlap, they cover much of the range above the atom, and the seams are checkable precisely because they overlap.

Below the atom this method is unavailable — matter there is discrete and there is no continuum to transform. That limitation is not an obstacle; it is the seam of §VII appearing again, this time in the instrumentation.

XI. Where this leaves the question

Does the vacuum resemble the weather?

As grammar: yes, and this is now measured rather than argued. Both are tension–release machines — potential accumulates, a threshold is crossed, discharge follows, recovery begins. Both have latent coherence below the threshold that noise excites rather than creates. Both carry the two-regime fluctuation signature, with the break at their own characteristic scale, as §IX shows for two layers with a calibrated null.

As identity — the same machine at a different layer — the claim has changed shape over the course of this argument, and for the better. It is no longer a single recursion with one spacing. It is two regimes joined at the atom: below, discrete scales derived from ratios of constants, of which one rung is now computed exactly as $(m_p/m_e) \times (1/\alpha)$; above, continuous scales set by dynamic balances, of which the healing length is the type specimen.

That is a weaker claim than the one this essay set out with, and a considerably more useful one, because both halves are now decidable and the two questions have different answers waiting in different places.

Is the climate a layer in that system? Yes, and it is the layer where the machine runs in full view. Its external drive is known, its records run to a century and a half, its internal wave dynamics are resolved by satellite and buoy array. The climate did not prove productive because it is more fundamental than anything else. It proved productive because it is the best-instrumented specimen of a machine that is otherwise mostly inferred.

The vacuum is not one more layer beside it. It is what is left when every layer of organisation has been stripped away — and its peculiarity, the thing that makes it the zero layer rather than merely the bottom one, is that it has exactly one knob where the climate has thousands.

There is one last convergence, and it is the reason this essay ends where it does rather than with a list of open problems.

A knotted electromagnetic field is not stable in free space. Solutions exist — the Hopf–Rañada fields carry genuine linked and knotted field lines — but they propagate and disperse; they are not standing bound states. What stabilises a knot is the nonlinear term. Without nonlinearity, no stable knot.

And a medium has no healing length without nonlinearity either. The competition that fixes ξ is precisely between a linear bending cost and a nonlinear interaction cost; remove the interaction and the length diverges and the medium becomes scale-free.

So the two halves of the argument — the knots below the atom and the dynamic balances above it — do not require two different things of the vacuum. They require the same single property, approached from opposite ends: **that the vacuum is nonlinear.**

Which is the only assumption the structure still rests on, and reduces a long question to a short one. *What fixes α, and what makes the medium nonlinear?* — and on the reading pursued here those may not be two questions.

Everything else about the structure of the vacuum is either a restatement of that, or invisible from where we stand.

Annotated References

The machine at the climate layer

Zebiak, S. E. & Cane, M. A. (1987). "A model El Niño–Southern Oscillation." *Monthly Weather Review* 115, 2262–2278. The minimal coupled ocean–atmosphere model, and the source of the single-dial structure this essay leans on. A reduced-gravity ocean beneath a fixed-depth mixed layer, coupled to a Gill–Matsuno atmosphere, with coupling strength μ as the control parameter. Everything downstream — the Hopf bifurcation, the noise-excited subcritical regime, the devil's terrace — is analysed in this model's parameter space.

Jin, F.-F. (1997). "An equatorial ocean recharge paradigm for ENSO." *Journal of the Atmospheric Sciences* 54, 811–829. The recharge oscillator. Explains the four-to-seven-year period through basin-scale heat content storage and discharge rather than through wave transit times alone. This is the tension–release cycle in its clearest published form and the direct climate-layer analogue of a vacuum that accumulates and releases stress.

Suarez, M. J. & Schopf, P. S. (1988) and **Battisti, D. S. & Hirst, A. C. (1989).** The delayed-oscillator mechanism: Kelvin waves propagating east, Rossby waves west, reflection at the western boundary returning a signal that kills the anomaly that launched it. Historically first, and instructive precisely because the wave timescales *fail* to reproduce the observed period — which is what motivated the recharge picture. Worth reading for how a mechanism can be real and still not be the explanation.

Matsuno, T. (1966) and **Gill, A. E. (1980).** The equatorial atmospheric response used inside Zebiak–Cane. Kelvin and Rossby waves on the equatorial beta-plane; the reason equatorial dynamics has a preferred direction and a finite signal-crossing time, which is what converts feedback into oscillation rather than runaway.

Jin, F.-F., Neelin, J. D. & Ghil, M. (1994, 1996). ENSO as a quasi-periodic route to chaos; the "devil's terrace." Where the fractal actually is. The interaction of ENSO's internal frequency with the seasonal cycle produces frequency-locked regions — Arnold tongues — separated by chaotic bands, and along a section at fixed coupling the locking ratios form a devil's staircase. Read alongside Bak.

Bak, P., Bohr, T. & Jensen, M. H. (1985). Circle maps, mode locking, and the transition to chaos. The abstract machinery: the circle map, the winding number, the tongue structure in the drive-frequency/nonlinearity plane, and the overlap of tongues above the critical line. This is the reason the same picture appears in the Pacific, in the asteroid belt, in cardiac arrhythmia, and in coupled oscillators generally.

Wisdom, J. (1982). “The origin of the Kirkwood gaps.” The asteroid-belt gaps as resonance overlap. Included here because it is the cleanest demonstration that the locking structure is not an analogy between the climate and the solar system — it is the same equation at two layers.

Scale, self-similarity, and how to tell strong claims from weak ones

Sornette, D. (1998). “Discrete scale invariance and complex dimensions.” *Physics Reports* 297, 239–270. The formal backbone for the layer claim. Distinguishes ordinary (continuous) scale invariance, which is weak and nearly unfalsifiable, from discrete scale invariance, in which self-similarity holds only at specific magnification ratios and leaves log-periodic corrections with complex critical exponents. If layers exist, this is the mathematics that says how far apart they are and how to detect the spacing.

Zhou, W.-X. & Sornette, D. (2002). *Physica D* 165, 94–125; and the follow-up correlation and spectral analyses. Direct detection of log-periodic corrections in the energy dissipation of high-Reynolds-number three-dimensional turbulence, with more than twenty harmonics of a single logarithmic frequency and a preferred scaling ratio near 2. The strongest existing evidence that discrete scale invariance is a physical phenomenon rather than a curve-fitting artefact.

Huang, Y., Johansen, A., Lee, M. W., Saleur, H. & Sornette, D. (2000). “Artifactual log-periodicity in finite-size data.” *Journal of Geophysical Research* 105, 25451–25471. The necessary companion to the above, and the reason the test in §VII carries its own null hypothesis. The same group that proposed the method published the analysis of when it produces spurious results. Any cross-scale claim that omits this control will not survive contact with a competent reader.

Sornette, D. (2009). “Dragon-kings, black swans and the prediction of crises.” The distinction between a black swan — self-organized criticality, power-law tail, the last grain on the sandpile — and a dragon-king, in which synchronization forces heterogeneous strongly-coupled units into a single rhythm, positive feedback drives the system to an extreme, and it collapses. Note that the mechanism is agnostic about content: it locates the runaway in the coupling between observers or agents, not in what is being observed. That makes it a sharp diagnostic and a blunt one at the same time.

Nottale, L. (2011). *Scale Relativity and Fractal Space-Time*. Imperial College Press. The most developed attempt to make scale itself a relativistic variable, with a non-differentiable fractal spacetime as the geometric substrate. Applications range from the QCD coupling and the cosmological constant to exoplanet orbital distances, Kuiper-belt object spacing, and stellar cycle durations. Whatever one concludes about the derivations, this is the reference case for "one scale law across many layers" pursued with quantitative predictions rather than resemblance.

Maeder, A. & Gueorguiev, V. G. (2017–2024). The Scale Invariant Vacuum paradigm; see *Universe* 8 (2022) 213 and *Symmetry* 16 (2024) 657 for reviews. Weyl integrable geometry with a gauge scalar field, applied from inflation and Big Bang nucleosynthesis to galactic dynamics, MOND, dwarf spheroidals, and lunar recession. Included because it is the closest thing in mainstream astrophysics to the position that the vacuum's scale properties, rather than its contents, are the operative variable — and because Gueorguiev works through the Ronin Institute and actively collaborates outside the university system.

Beck, C. (2002). *Spatio-temporal Chaos and Vacuum Fluctuations of Quantized Fields*. World Scientific. Deterministic chaotic models of vacuum fluctuations at small scale, using coupled map lattices ("chaotic strings") as the noise source for stochastic quantization in the Parisi–Wu scheme, with numerical evidence that the vacuum energy is extremized at the observed standard-model parameter values. Directly relevant to §IV: an explicit attempt to derive coupling constants from a dynamical system rather than to postulate them.

The vacuum as a medium

Sakharov, A. D. (1967). "Vacuum quantum fluctuations in curved space and the theory of gravitation." *Doklady Akademii Nauk SSSR* 177, 70–71. Induced gravity: the Einstein–Hilbert action as the elastic response of the vacuum to deformation, gravity as a medium property rather than a fundamental interaction. Four pages, and the origin of the entire vacuum-elasticity line of thought.

Volovik, G. E. (2003). *The Universe in a Helium Droplet*. Oxford University Press. The sustained development of the same idea with a concrete laboratory system. Superfluid ^{3}He-A generates emergent chiral fermions, gauge fields, and an effective metric as collective modes. The book's decisive lesson for this essay is epistemic rather than physical: observers made of a medium's low-energy excitations are structurally unable to see the medium's own architecture. That is the general form of the argument made specifically in §IV.

Williamson, J. G. & van der Mark, M. B. (1997). "Is the electron a photon with toroidal topology?" *Annales de la Fondation Louis de Broglie* 22, 133. Charge as a topological winding number of confined light; mass as trapped circulation. The proposal that makes $\alpha = e^2 Z_0/2h$ read as a statement about the relation between a medium and its knots rather than as a numerical accident.

van der Mark, M. B. (2015). "On the nature of 'stuff' and the hierarchy of forces." *Proceedings of SPIE* 9570. An energy-balance argument tracking the internal kinetic energy fraction from the Earth–Sun system through hydrogen and the deuteron to the proton, concluding that granularity must terminate where that fraction approaches unity. The consequence — a continuous, topologically electromagnetic substrate with the proton as the smallest stable knot — is the cleanest available statement of what "the vacuum is the zero layer" would mean physically.

BIPM (2019). *The International System of Units,* 9th edition; and CODATA recommended values. The reference for the constants argument. Since the 2019 redefinition, c, h, and e are exact by definition, with the consequence that μ_0 carries an experimental uncertainty inherited from α. Worth consulting directly rather than through secondary sources, because the direction of the dependency is what makes the single-knob argument work and it is easy to state backwards.

Cycles across domains

Dewey, E. R. & Dakin, E. F. (1947). *Cycles: The Science of Prediction*. Henry Holt. The foundational compilation. Dewey's conclusions after collating several thousand cycle reports: everything that can be made into a time series has cycles; the same periods recur across unrelated fields; those shared periods are phase-synchronized; and the ratios between commonly reported periods are dominated by small integers, especially 2 and 3. Whatever one makes of the interpretation, the compilation itself is a substantial empirical object and it has never been properly explained.

Tomes, R. *Harmonics Theory*. https://ray.tomes.biz/ — see in particular *The Physics and Mathematics of the Harmonics Theory* (https://ray.tomes.biz/maths.html) and *The Harmonics Theory: the story of how it came about* (https://ray.tomes.biz/story.htm). Cycles Research Institute, Auckland. The central reference for §V and §VI, and the most complete existing attempt to derive the layer structure of the universe from a single recursive principle. One axiom — the universe consists of a standing wave which develops harmonically related standing waves, and each of these does the same — with the recursion following from the nonlinearity of Maxwell's equations under general relativity rather than being postulated. The strength of each scale is determined by the number of routes along which it can be reached, which reproduces the intervals of just intonation as the strong scales of the medium, and thereby gives the music of the spheres a mechanism rather than a metaphor.

The calculations run to harmonics beyond 10^50, comparing predicted scale ratios against the observed hierarchy from the Hubble scale down to nucleons, with an average predicted spacing of 10^4.56 against roughly 10^4.5 observed. The site also supplies the computed tables as downloadable CSV, including the main-line harmonics, which makes the work independently checkable rather than merely assertable — an unusual courtesy.

Three things distinguish this from the general run of cross-scale theories. It was arrived at from data, not from theory: economic cycles in the New Zealand economy at exact fractions of 35.6 years, then musical ratios in forty-four years of weekly corn prices, and only afterwards the discovery that Dewey had independently found the same periods. It is top-down, forming structure from the largest scale inward, which is the correct orientation for a picture in which the vacuum is the residue of stripping organization. And its author states plainly where it does not yet close — the ratios 11, 17 and 19 appearing in the right order at the wrong spacings — rather than absorbing the discrepancy.

The cosmological consequences Tomes draws (an oscillating universe much older and larger than the Hubble scale, no Big Bang, and a fundamental period near 1.48 × 10^23 years) are separable from the axiom and the layer spacing, and a reader who rejects the former need not reject the latter.

Bejan, A. Constructal law. Flow systems evolve toward configurations that provide easier access to their currents, producing the same branching tree at every scale — lungs, river deltas, vasculature, cities. The complementary fractal to Tomes's: repetition of a *shape* across size scales rather than repetition of a *locking structure* across ratio scales. Both may be operating; they are not the same claim and should not be merged.

Gunderson, L. H. & Holling, C. S., eds. (2002). *Panarchy: Understanding Transformations in Human and Natural Systems*. Island Press. The adaptive cycle — growth, conservation, release, reorganization — nested across scales, with faster small cycles and slower large ones coupled by "revolt" and "remember" transitions. The ecological statement of the same tension–release machine, and the clearest existing treatment of how layers of one nested system interact rather than merely resemble each other.

Method for the test

Hurst, H. E. (1951). "Long-term storage capacity of reservoirs." *Transactions of the American Society of Civil Engineers* 116, 770–808. The original rescaled-range analysis, developed on Nile flood records. The exponent that carries the test in §VII, discovered in exactly the kind of long geophysical series the test proposes to use.

Mandelbrot, B. B. & Van Ness, J. W. (1968). "Fractional Brownian motions, fractional noises and applications." *SIAM Review* 10, 422–437. The formal model behind H: fractional Brownian motion, long-range dependence, and the relation between the Hurst exponent, the spectral slope, and the fractal dimension of a trace. The mathematical object the cross-layer comparison is comparing.

Kolmogorov, A. N. (1941). The −5/3 inertial-range spectrum. The reference exponent for a fully developed turbulent cascade, and the benchmark against which solar wind and atmospheric spectra are read.

Bruno, R. & Carbone, V. "The solar wind as a turbulence laboratory." *Living Reviews in Solar Physics* (updated periodically). The best single review of the layer that matters most for the test: a nearly collisionless magnetized medium whose fluctuation spectrum is measured directly over several decades of scale, with a Kolmogorov-like inertial range and a spectral break at ion scales. If any accessible system is close to "the electromagnetic medium observed on its own terms," this is it.

On the shape of the argument itself

Taleb, N. N. (2012). *Antifragile: Things That Gain from Disorder*. The systematic treatment of gaining from volatility rather than merely surviving it, and — read against §I — a useful corrective to the assumption that noise is a nuisance to be filtered out. In the subcritical Zebiak–Cane regime the noise is not degrading the signal; it is the only reason there is one.

Wheeler, J. A. & Feynman, R. P. (1945, 1949). Absorber theory. Advanced and retarded solutions taken equally seriously, with the standing wave between emitter and absorber as the physical object. Included because the one-electron intuition and the confined-circulation reading of charge both descend from taking this seriously, and because Wheeler framed it in particle language at a moment when the wave reading had not yet been absorbed — which is why the idea is usually remembered as a curiosity rather than as a program.

Van der Pol, B. (1927). "Frequency demultiplication." *Nature* 120, 363–364. Two pages. Relaxation oscillations in a driven triode circuit, the frequency-locking staircase found by turning a knob, and — recorded almost in passing — the irregular noise heard in a telephone receiver just before each jump. The first observation of deterministic chaos, made by someone with his hands on the apparatus. A reminder of where in the process the discoveries in this line have actually tended to occur.

Knots as solitons: why a winding number can be stable

Faddeev, L. & Niemi, A. J. (1997). "Stable knot-like structures in classical field theory." *Nature* 387, 58–61. The founding paper of the modern knot-soliton programme, and explicitly presented as the first genuine realisation of Kelvin's 1867 proposal that atoms are knotted vortices — an idea that waited 130 years for a mechanism because nobody could tie a knot in a field. The solitons are classified by an integer Hopf charge, which is a linking number rather than a degree. This is the machinery the "why four" question of §X requires.

Battye, R. A. & Sutcliffe, P. M. (1998). Knots as stable soliton solutions; and **Sutcliffe, P. M. (2007).** "Knots in the Skyrme–Faddeev model." *Proc. R. Soc. A*. The numerical spectrum. Battye and Sutcliffe reached Hopf charge eight and identified the first true knot — the trefoil — as the minimum-energy configuration at charge seven. Sutcliffe extended to charge sixteen and catalogued a range of torus knots, with the configurations changing character as the charge rises: buckling at three, linked pairs at five and six, the trefoil at seven. The fact that matters most for §X is the boundary: the minimum-energy state remains a twisted torus — a single closed loop — only up to charge four, and from five upward it is a link or a knot. Later minimisation work reports revised ground states at charges four and five and a symmetry-breaking deformation at six, improving the fit to the Vakulenko–Kapitanskii bound $E \sim |Q|^{3/4}$, so the charge-four shape is still under revision. Note also what these papers do not fix: the model determines shapes and energy ratios but not absolute size, which follows from the ratio of the two Lagrangian terms and must come from outside.

Rañada, A. F. (1989–1992); and reviews of knotted electromagnetism, e.g. *Physics Reports* (2016). Knotted solutions of Maxwell's equations in vacuum — the Hopf–Rañada fields, with genuinely linked and knotted field lines. Essential to read alongside Faddeev–Niemi rather than instead of it, because these solutions propagate rather than standing still: free-space electromagnetism gives knot topology but not knot stability. The contrast is what forces the nonlinearity conclusion of §XI. The same literature supplies the selection rule used in §X: torus-knot magnetic configurations require poloidal winding number unity, while the toroidal winding may be any integer.

Kleckner, D. & Irvine, W. T. M. (2013). "Creation and dynamics of knotted vortices." *Nature Physics* 9, 253–258; and Irvine's subsequent work on knotted light fields. Kelvin's proposal finally performed rather than argued: isolated trefoil vortex knots created in water using rapidly accelerated 3D-printed hydrofoils, and their three-dimensional evolution tracked by high-speed tomography. The result matters as much as the method — the knots are not stable. They stretch themselves and then reconnect, passing from knotted to unknotted. Frozen topology that unties itself, with a measurable mechanism. The nearest thing to experimental work on the §VI–§X questions that currently exists.

Condensates and internally generated length scales

Pitaevskii, L. & Stringari, S. (2016). *Bose–Einstein Condensation and Superfluidity,* 2nd ed. Oxford University Press. The standard reference for the healing length and the Bogoliubov spectrum used throughout §VI and §VII: the competition between kinetic bending cost and interaction cost, the resulting single internal length, and the crossover from collective (linear dispersion) to single-particle (quadratic) behaviour at that length. The technical content behind the claim that a condensate is not scale-free.

Sikivie, P. & Yang, Q. (2009). "Bose–Einstein condensation of dark matter axions." *Phys. Rev. Lett.* 103, 111301; and **Marsh, D. J. E. (2016).** "Axion cosmology." *Physics Reports* 643, 1–79. The demonstration that a condensate can be treated quantitatively on cosmic scales, and the standard review of the resulting phenomenology. In ultralight-boson models the de Broglie wavelength reaches kiloparsec scales and the condensate produces solitonic cores at the centres of haloes — a medium's internal length made visible as macroscopic structure, and measurable in rotation curves. Cited here for the mechanism rather than for the dark-matter claim; the Gross–Pitaevskii–Poisson equations are indifferent to what the constituent is.

Fits of the Gross–Pitaevskii–Poisson formalism to dwarf galaxies, e.g. the 2026 analysis of DDO 168 giving a solitonic core radius of 2.4 kpc. The worked example behind §VII. Whatever one concludes about ultralight bosons, this is what "a medium whose own parameters fix one observable length" looks like when it is actually fitted to data.

The scale-invariant vacuum, and why it is not this

Gueorguiev, V. G. & Maeder, A. (2024). "The Scale-Invariant Vacuum paradigm: main results and current progress review (Part II)." *Symmetry* 16, 657; and Maeder, A. (2017), *ApJ* 834, 194. The most developed contemporary programme in which scale is the vacuum's operative variable, built on Weyl integrable geometry, with a single scalar λ fixed by a gauge condition to $\lambda = t_0/t$ — one degree of freedom, and determined rather than free. Two features make it worth reading closely here. Its reparametrisation argument reads the extra κv term as an artefact of describing a process in the observer's time rather than its own, which turns a metaphysical claim about time into a term in an equation. And its lunar-recession result — 0.92 cm/yr closing the gap between the 3.83 cm/yr measured by laser ranging and the 2.85 cm/yr implied by the historical slowing of the Earth — is the one place where this class of theory touches a solar-system measurement.

Two structural features also make it incompatible with the picture drawn here, and the incompatibility is worth stating plainly. Its λ is forbidden to depend on position — homogeneity and isotropy force $\partial_i\lambda = 0$ — so structure, which is by definition what varies from place to place, cannot exist in the framework as formulated. And its scale invariance is the continuous kind: no preferred scale, no privileged magnification. The authors themselves name the spatial dependence of λ as an open direction in their conclusions.

There is a further consequence they do not draw. Their own equations make the effects of scale invariance grow as matter density falls, vanishing entirely at and above critical density. An Earth-based laboratory is therefore, on their own account, the worst possible place to look — a null result there is predicted rather than informative. There exists a whole class of theories for which terrestrial experiment is structurally blind.

Where the constants are actually measured

Budker, D. and the GNOME collaboration. The Global Network of Optical Magnetometers for Exotic physics searches. A worldwide network of synchronised magnetometers built to detect transient, correlated signals as the Earth passes through a region where field values differ. If the vacuum has local structure, this is the class of instrument that would see it — not as a constant offset, which is undetectable, but as a passing edge.

Optical clock networks, e.g. the European fibre-linked comparison constraining $|\delta\alpha/\alpha| < 5 \times 10^{-17}$ at $\sim 10^3$ s. The direct version of the operational corollary in §IV. Since any local vacuum structure can express itself only as a local variation of α, a network searching for transient α variations as spatially localised structures pass through is searching for exactly the thing this essay says is the only observable. Published limits exist; what does not exist is a model saying what should be passing.

Webb, J. K. and collaborators; and the competing systematics analyses. Three decades of measuring α in quasar absorption spectra, including the reported spatial dipole, together with the sustained internal criticism of those measurements. Worth reading with the criticism rather than around it: the argument over systematics is where the real epistemics of “does a laboratory result generalise” is being conducted.

Scaling regimes in geophysical data

Lovejoy, S. *The Weather and Climate: Emergent Laws and Multifractal Cascades* (with D. Schertzer, Cambridge, 2013), and subsequent papers. Forty years of measuring scaling exponents in atmospheric and climate data, and the closest existing framework to the measurement reported in §IX. Lovejoy identifies a transition from a weather regime to a “macroweather” regime in which successive fluctuations cancel, and a further transition to a climate regime above roughly thirty years. The exponent of 0.58 measured here beyond the ENSO break — fluctuations averaging themselves out — is his macroweather regime, and his maps of transition timescales single out the El Niño region off Peru as distinctive. Anyone continuing the §IX measurement should start here rather than repeat it.