Ψ

PIU-Ψ

Universal Integrity Principle

One single thing

What is the cosmos? What is the Pleno? How does it operate?

A research program in fundamental physics

An explanation to grasp the idea, almost without mathematics

Manuel Alberto Celedón Mejía • Independent researcher
ORCID 0009-0002-8090-2364 • Medellín, Colombia • piuniversal.com

Popular-science version V5.1 • August 2, 2026
Zenodo DOI: 10.5281/zenodo.21724390

The underlying question

Imagine someone told you that the gravity holding you in your chair, the light that lets you read this, the particles you are made of, and the great mysteries of the cosmos —dark matter, dark energy— are all the same thing, seen in different ways. Not separate forces, not different ingredients: a single substance, behaving in diverse forms.

To understand why that idea is so powerful, you first have to know about an enormous problem physics has carried for a century. Today we describe nature with two great theories that work wonderfully... but that do not speak to each other. On one side, the theory of the very large: general relativity, which explains gravity, the stars, the galaxies, and the shape of the entire universe, with astonishing precision. On the other, the theory of the very small: quantum mechanics, which governs particles, atoms, and the subatomic world, just as precise. Each reigns in its own domain; together, they contradict each other. When one tries to use the rules of the quantum world to describe gravity, the mathematics fills with meaningless infinities and everything collapses. The finest minds of the past century have tried to bridge that gap —they call it "quantum gravity," and it is considered the greatest unsolved challenge in fundamental physics— without success. It is like having two perfect maps of the same city that, overlaid, do not fit.

Faced with that wall, PIU-Ψ takes a bold and different step, and its most original trait lies in what that step consists of. Almost every previous attempt accepts that there are two theories and looks for a way to *unite* them. PIU-Ψ inverts the question: **what if they were never two?** What if the large and the small are not two worlds to reconcile by force, but two faces of a *single* physical substance? That substance is called the Pleno, and from it —the program proposes— everything else emerges.

Three clarifications, from the outset, to situate what follows:

- **It is physics, not philosophy.** PIU-Ψ is not a metaphor or a poetic vision of reality. It is a **research program in fundamental physics** —the branch that seeks the ultimate laws of nature, the same one in which Newton, Einstein, or Schrödinger worked—, built, like any serious physics, on **high-level mathematics**: it starts from precise principles, derives its results the way theorems are proven, and states exactly what is proven and what is not. Everything that follows —the crystal, the tension, the spin, time— are images in words of objects that, underneath, have their exact equation.

- **It does not come to replace, but to explain.** PIU-Ψ does not aim to repeal relativity or quantum mechanics. On the contrary: it sets out to *derive* them —to show that both spring from a single root— and thereby **integrate them from a common base**, not supplant them. The laws that already work will keep working; what the program adds is the trunk from which they hang.

- **It is a new approach.** Its originality lies not in proposing yet another "theory of everything," but in the *change of question*: instead of forcing two pieces to fit, it suspects they were always one. That turn is what lets it sidestep the infinities that stalled the others.

This document tells, with everyday images and almost without mathematics, what the Pleno is and how it operates. The interleaved **blue boxes** give the technical version of each stretch —the equation, the number, and its status—; the thread follows just as well without them.

1. From a single equation, the laws of the universe

The starting point of PIU-Ψ is that same question —what if everything, gravity, light, matter, the cosmos, were at bottom the same?— carried to its ultimate consequences. Many have felt, at some point, that suspicion of unity beneath the diversity of all there is. What is distinctive about PIU-Ψ is not to settle for the idea, but to submit it to the most demanding tribunal that exists: that of mathematics. The reasoning is simple and strict: if everything truly is a single substance, then, starting from a precise mathematical description of that substance, one must be able to derive —rigorously, step by step, without cheating— the known laws of physics. And if one could not, the idea would be refuted.

That is exactly what the program does. PIU-Ψ proposes that the entire universe is a single field —the Pleno—, and writes it with the same kind of mathematical object physics already uses for the smallest things, now extended to everything that exists. The symbol Ψ alone already hints at the great wager: that the quantum is not a separate corner of reality, but the very texture of the single substance that forms it all. That field is written like this:

$$\Psi = S \cdot e^{i\theta},$$

and no mathematics is needed to read the essential part: the formula says that at each point the Pleno has only *two* qualities —S, how much it is tensed, and θ, which way it spins—. Nothing more. From those two quantities, and how they change from one point to the next, everything else will spring.

That expression describes *what* the Pleno is. To know *how* it behaves, the theory adds its "fundamental equation," known in PIU-Ψ simply as F1. There is no need to understand it: it is enough to know what it is, a single compact recipe that fixes the dynamics of the Pleno. That equation is not chosen at random nor adjusted along the way: it is built from a few starting principles —about half a dozen precise statements about how the Pleno is, the Axioms of PIU-Ψ— and, once fixed, it is closed, immovable. From there, everything else follows by strict protocol and mathematical deduction, without introducing new ingredients. The theory derives —one after another, with the rigor of a theorem— the laws physics discovered separately over the centuries:

1. Newton's gravitational constant, and the structure of Einstein's equations.
2. Planck's quantum constant and Heisenberg's uncertainty relation, with its exact floor.
3. Maxwell's equations of electromagnetism, and Yang-Mills' of the nuclear forces.

4. Dirac's equation of the electron, with its spin and relativistic structure.

5. The Chandrasekhar limit that decides the death of a star.

6. The Stefan-Boltzmann law and Bose-Einstein statistics, obtained by counting the configurations of the medium itself.

7. And the proportions of the cosmos —how much dark matter and dark energy there is— in figures that fall within the satellites' own margin of error.

Each of those names —Newton, Einstein, Maxwell, Yang, Heisenberg, Schrödinger, Dirac, Lorentz, Chandrasekhar, Bose— represents a lifetime of work, sometimes of whole generations, and each law was found separately, in different eras, to explain phenomena that seemed to have nothing in common: the gravity of the planets, the magnetism of a magnet, the behavior of an electron, the death of a star, the glow of a hot body. Different chapters of physics, in different books. What PIU-Ψ proposes is that all those chapters are, at bottom, the same story —that of the Pleno—, and can be written with the same pen, from the same first line. That is its ambition and most distinctive trait: not to have explained a new phenomenon, but to have found the common trunk from which branches we took for separate trees grow.

The starting point, in its exact form

Six axioms fix what the medium is: **A1**, a single, finite, conserved energy, with no vacuum; **A2**, a lower wall preventing the medium from vanishing; **A3**, causality at finite speed c; **A4**, a grain of size ℓ_P; **A5**, an internal phase; **A6**, upper saturation. Reading A1 region by region gives the density of the medium, ρ_P. And writing conservation as a balance between two orthogonal branches —what the medium retains as structure and what it propagates as change— forces the form of the field: $\Psi = S\,e^{i\theta}$, where S is how tensed the medium is and θ which way its phase points. The complex field is not a computational convenience but the minimal expression of that balance.

Applying the Lagrangian procedure to that field under the six axioms produces F1:

$$\boxed{\mathcal{L}_{F1} = \tfrac{1}{2}\rho_P c^2 \ell_P^2\, \partial_\mu \Psi^* \partial^\mu \Psi - V(S) - \frac{\rho_P c^2 \ell_P^4}{40}(\Box S)^2 - V_q^{\text{dens}}(S)}$$

Four terms, each with a structural role: the **kinetic** one propagates the field, carrying the density of the medium ρ_P and its grain ℓ_P in the coefficient; the **potential** $V(S)$ holds the stable configurations; the **flexural** term penalizes curvature of the modulus; and the **Klauder** term, V_q^{dens}, diverges as $S \to 0$ and keeps the value zero inaccessible. Stretching a sheet and bending it are two different ways of deforming it, and a material can resist one strongly and the other hardly at all: a sheet of steel bends by hand but there is no stretching it. The flexural term is the Pleno's resistance *to being bent* —to its tension changing in an ever tighter way from one point to the next— and its cost grows with how tight that curvature is. What it penalizes, therefore, are variations of a length comparable to the grain: the material itself objects to anything finer than itself occurring. None of the four can be removed without sacrificing functionality, and none was added out of a need to fit. And F1 is written with **three primary quantities**: the density of the medium ρ_P, its maximum rate of change c, and its grain ℓ_P. Everything else appearing in the program —G, $\hbar$, the scale of the potential, the two walls— follows from those three plus the geometry of the lattice. The Euler-Lagrange equations on F1 return the modulus equation, the phase equation, and an exactly conserved current.

What reaching this point means. Everything else in the program —gravity, the quantum, the cosmos— is *evaluating* this object, not extending it. And that carries a hard consequence: a theory with a single dynamical object cannot accommodate a disagreement by adding a term to it. Every result arrives with these same coefficients or it does not arrive. **[DERIVED]**

2. What "derive" means

That word —derive— will appear on every page from here on, and the whole weight of what the program claims depends on what it means.

There are two ways to arrive at a formula in physics, and the distance between them is the distance between describing and explaining. The first is to **fit**: take a measurement, propose a formula with a few knobs, and turn the knobs until the formula passes through the data. It is legitimate and has built much of our knowledge, but it has a known limit: a formula with enough knobs always ends up succeeding, so success by itself proves little. The second is to **derive**: start from a fixed point and advance by steps in which each one follows from the previous without introducing anything new along the way, just as a theorem is proved. Here there is nothing to turn, and so the result can come out wrong. That it comes out right does prove something.

And how does one derive anything from an equation like F1? A Lagrangian is an energy balance: it says how much energy the system has in each configuration it could adopt. There is a mathematical procedure —old, standard, the same one all of theoretical physics uses— that takes that balance and returns the equations of motion: how the system evolves, what stable shapes it admits, what quantities are conserved in it. That procedure admits no creativity. Applied to the same Lagrangian, two people obtain the same equation, just as two people dividing the same numbers obtain the same quotient.

From that general equation to a known law there is one more step: looking at it in a particular **regime**. The same equation behaves differently depending on whether it is observed at low speeds or close to the speed of light, at enormous distances or at the size of the grain, with configurations disordered or all aligned. Restricting the view to a regime simplifies the equation, and what remains is often a formula physics already had. That is an *effective limit*, and it is the mechanism by which a deeper theory contains its predecessors instead of repealing them: Newton's mechanics did not cease to be true when Einstein arrived, it remained as the limit of relativity for low speeds and weak fields.

There lies the meaning of so many results coming out of a single F1. Each time a regime is chosen and the computation is done, the theory has no margin: it returns the number it returns. If F1 gave a gravity too strong, or a radiation with the wrong coefficient, or a proportion of dark matter that did not match, there would be nothing to retouch —F1 is the same object in every case, and a knob that fixed one result would wreck the others—. That laws obtained in unrelated regimes —gravity and electric charge, the glow of a hot body and the composition of the cosmos— all come out of the same object with nothing adjusted is precisely what a fit cannot imitate.

The procedure, and what an effective limit is

The step from a Lagrangian to the equations of the system is taken by the Euler-Lagrange equation, which for a field Ψ reads

$$\partial_\mu\left(\frac{\partial\mathcal{L}}{\partial(\partial_\mu\Psi)}\right) - \frac{\partial\mathcal{L}}{\partial\Psi} = 0.$$

The first term measures how the Lagrangian changes when the *variation* of the field from one point to the next changes; the second, how it changes when the field itself changes. Setting them equal to zero selects, among all imaginable histories of the system, the one for which the action does not vary to first order. It is a closed procedure: given the Lagrangian, the equation is fixed.

The power of that tool shows in what it already returns in known physics. Applied to the Lagrangian of a classical particle it reproduces Newton's laws; to that of the electromagnetic field, Maxwell's equations; to the Einstein-Hilbert action, the field equations of general relativity. Applied to F1, it returns the dynamical equation of the Pleno —one for the modulus, another for the phase— and the conserved current accompanying the phase symmetry, by Noether's theorem.

An **effective limit** is what remains of that equation when it is restricted to a regime. Formally it is an expansion in a small parameter: in powers of v/c when speeds are low, in powers of $k\ell_P$ when wavelengths are large compared with the grain, or by averaging over configurations when collective behavior is what matters. Higher-order terms are dropped and a simpler equation appears, with its domain of validity explicit.

What this means for the program. All the content is already inside F1 from the first moment; the work consists in extracting it, not inventing it. That is why the program's results are not a collection of independent proposals to be accepted or discarded one by one, but a tree of chained consequences: each hangs from the previous, and removing one link leaves everything after it unsupported. **[DERIVED]**

3. A single substance that fills everything

Imagine that the universe is not made of loose things floating in a void, but is a single continuous material that fills absolutely everything, without gaps. What we call "empty space" would not be empty: it would be that material at rest.

A good image is that of an elastic medium —think of an infinite gelatin, or a transparent crystal extending in all directions—. And everything that exists would be a shape that medium takes. Just as waves, whirlpools, and currents of the sea are not things distinct from water but forms of water, particles, light, and gravity would not be things distinct from the Pleno, but forms of the Pleno. A tension here, a spin there, a wave running through it. That is the central wager of PIU-Ψ, and everything else hangs from it: a single substance, and the entire universe as its infinite ways of moving.

4. How it is organized: a diamond-shaped crystal

A medium like this is not disordered mush. The Pleno organizes itself, structures itself internally like a crystal, with its units arranged in a regular, repeated pattern, just like the atoms in a gemstone.

And not just any pattern, but the same one a diamond has. In a diamond, each carbon atom bonds to four others, arranged toward the tips of a tetrahedron —picture four particles arranged so that they always form a triangular-based pyramid—. It is the most balanced and firm way to distribute the bonds in space, which is why diamond is the hardest natural material there is. The Pleno shares that architecture: each of its cells bonds to four others in that same tetrahedral arrangement. From that simple geometry —"four neighbors in a tetrahedron"— follow a great many of the numbers on the earlier list. The shape of the crystal is not an ornament: it is the source of its rules.

The diamond structure is not postulated, it is *derived*. It is not a convenient assumption slipped into the theory "just because." When one asks how a medium like the Pleno must organize itself to be stable and self-consistent, the mathematics point, unambiguously, to that tetrahedral arrangement —the same one nature chooses, not by chance, for the most resistant material there is—. And from that single geometric form follow, chained and also with proof, a multitude of the results the theory produces. The geometry of the crystal is not an ornament nor a choice: it is a derived consequence, and at the same time the root from which almost everything else hangs: the angle of the bonds and the number of neighbors fix concrete numbers that reappear throughout the edifice.

The lattice and the two pure ratios of the medium

F1's potential has a wall below and a wall above, but the form of the Lagrangian does not by itself say at which value of the modulus each one falls. That is fixed by the microstructure the medium adopts: on evaluating F1 one sees that it crystallizes into the diamond lattice, four neighbors per cell —the structure that minimizes F1's own flexural term, the one that costs least to bend, with a uniaxial discriminant excluding the rival stackings—. And from that lattice come the two positions.

The **floor** is computed by two independent routes. The first is purely geometric: average the discrete

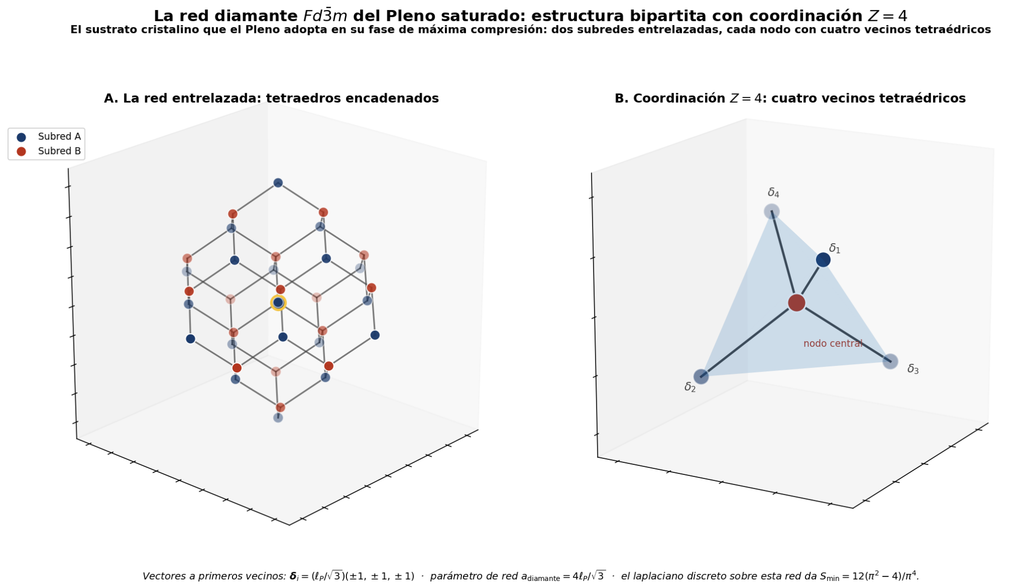

Figure 1: The internal structure of the Pleno: each cell bonded to four others in a tetrahedral arrangement, the same architecture as diamond. From this geometry come the two pure ratios of the medium.

Laplacian over every direction of the lattice and compare it with that of its even sublattice,

$$S_{\min} = \frac{\langle D_{\text{lat}} \rangle}{\langle D_{\text{lat}}^{\text{even}} \rangle} = \frac{12(\pi^2 - 4)}{\pi^4} = 0.7230871.$$

The second is dynamical and shares not one algebraic step with the first: find the minimum of the potential, solving $V'(S) = 0$. It gives 0.7230874. They agree to seven figures.

The **ceiling** comes from maximal sphere packing —Hales' theorem—: how far the diamond lattice can be compressed before saturating,

$$S_{\max}^{2,\text{geom}} = \frac{n_{\text{FCC}}^{\max}}{n_{\text{diamond}}} = \frac{8}{3}\sqrt{\frac{2}{3}} = 2.1773242,$$

that is, a maximum modulus $S_{\max} = 1.4756$. And the coefficient $1/40$ of F1's third term comes from that same geometry.

What reaching this point means. The agreement of the two routes was not guaranteed: they are different computations on different objects, and a discrepancy in the third figure would have meant that the Lagrangian and the lattice do not describe the same medium —and the program would have broken there—. That they agree is the check that they do. From here on, the numbers of the program are not chosen: they are read off a geometry which was itself not postulated but obtained by minimizing F1. **[DERIVED]**

5. The two qualities of each point: how much it tenses, and which way it spins

And how does one describe what happens in the Pleno? Surprisingly, with only two pieces of data at each point —the same two that already appeared in the formula $\Psi = S \cdot e^{i\theta}$—:

- **How much it is tensed or compressed** —this is the S of the formula—. Call it its *tension*. At rest, that tension has a background value; where there is "something" —a particle, a wave—, the tension rises or falls.

- **Which way it is spinning** —this is the θ of the formula—. The Pleno has, at each point, a kind of spin orientation, like the hand of a tiny clock. What matters is not the hand of a single point, but how the orientation changes from one point to the next: if they all point the same way, the medium is calm; if the orientation rotates from one point of the medium to the next, that is a wave traveling.

With these two qualities —tension and spin— almost everything is understood. From tension, particles and matter are born; from spin, light and the forces. The tension of the Pleno never reaches two extremes. It cannot stretch until it vanishes (it has a floor, never reaching zero) nor compress without end (it has a ceiling). The medium always keeps “some” substance and never collapses entirely.

Those two limits are not a precaution added by hand. They are the two walls of a well: the Pleno has an energy curve of its own, with a bottom where it rests and a wall on each side, and from that curve come both dark energy and the impossibility of the universe collapsing to a point.

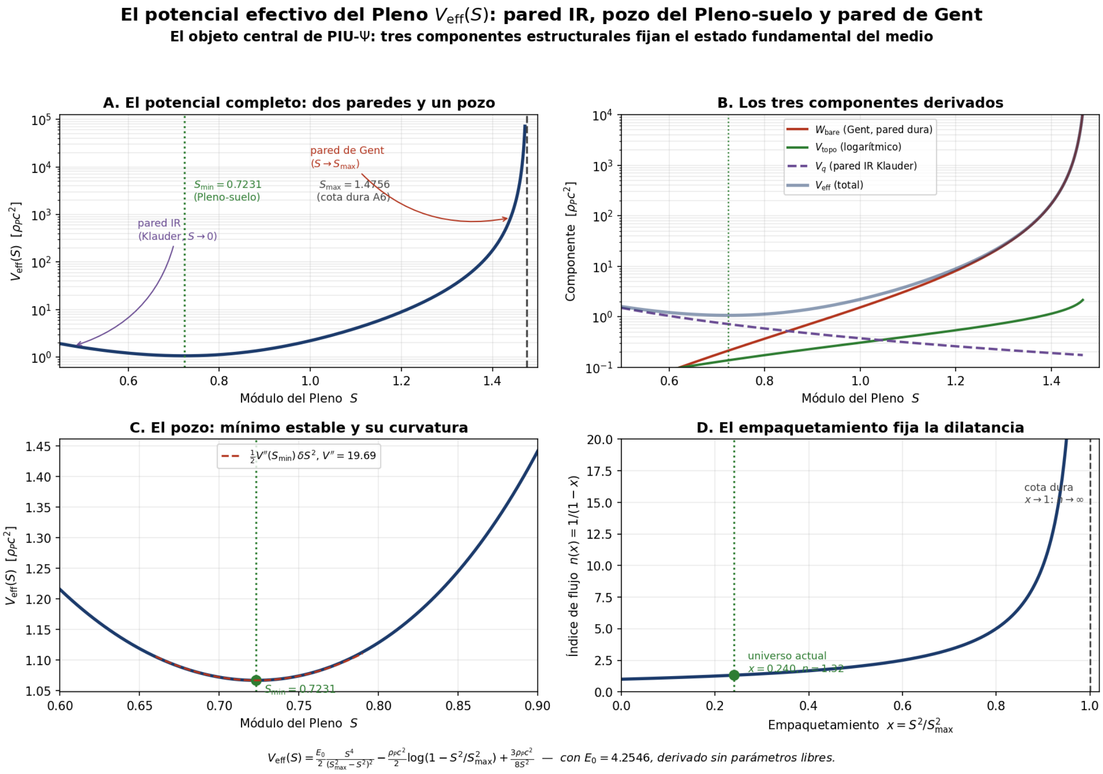

Figure 2: The effective potential of the Pleno: the lower wall preventing the medium from vanishing, the well where it rests today, and the hard wall preventing endless compression. All three pieces are derived, with no free parameters.

The well, and why the medium cannot leave it

The two bounds are not stipulations: they are the two walls of a curve F1 carries within. Adding the three components of the potential gives the effective potential $V_{\mathrm{eff}}(S)$, which says how much energy each configuration of the medium costs, and its shape is a well with a wall on each side.

The **lower wall** comes from the Klauder term, which diverges as the modulus tends to zero: vanishing would cost the medium infinite energy, so the absence of vacuum is not a prohibition imposed from outside but an energetic impossibility. The **upper wall** comes from the logarithmic Gent term, which diverges as S approaches its saturation value: compressing the lattice beyond maximal packing likewise costs infinite energy.

Between the two walls lies the **bottom**, and there the medium rests today. Its position is $S_{\min} = 0.7231$ —the same number the lattice gave— and the curvature of the well at that point is

$$\xi = \frac{V''_{\mathrm{eff}}(S_{\min})}{\rho_P c^2} \approx 19.69,$$

a dimensionless number measuring how strongly the medium resists being moved from its rest. In that ground configuration the Pleno holds 52% of its primary structural density: $\rho_{\text{ground}} = \rho_P\, S_{\min}^2 \approx 0.5229\, \rho_P$.

What reaching this point means. The potential is not chosen from a family of possible shapes so the numbers come out, which is the usual procedure when a model needs a potential: its three components are derived and its energy scale, $E_0 = 4.2546$, comes out in closed form in π and $\sqrt{6}$, with not a single free symbol. A well with no free parameters is what allows whatever is deduced from it to be a consequence and not a fit. **[DERIVED]**

6. A firm medium that, when it must, flows

The Pleno has two distinct ways of resisting deformation, and from each springs a famous consequence of physics. The first we can call its *firmness*; the second, its *resistance to abrupt change*. They are two different properties of the same medium, not one.

Firmness: how much the medium yields when you push it. A good way to measure how firm something is, is to see how much it yields when you press it. Sink your finger into a pillow: it sinks a lot. Rest the same finger, with the same force, on a wooden table: it barely yields. The pillow is soft; the table, firm. "Firm" means no more than that: that for a given push, it yields little.

Now, the Pleno —even at rest, calm, which is how it is today— is extraordinarily firm in that sense: when something pushes it, it yields very little. That the Pleno is firm does not mean it is tense or compressed, like a rope stretched to the point of snapping. On the contrary: today the medium is in its most *relaxed* state, calm, settled in its rest —the state the program identifies with dark energy—. It is the serene firmness of something deeply settled, not the tautness of something forced. From this firmness at rest will come, against all intuition, the weakness of gravity.

Resistance to abrupt change: why the speed of light exists. The second property is of another nature, and here one of the most elegant ideas appears. That firm medium normally behaves like an elastic solid: deform it a little, and it recovers its shape, like a spring. But it also has a character that depends on *haste*: the faster it is forced to deform, the more it resists, until it becomes infinitely rigid in the limit. It is not the same as its static firmness —it is a resistance that grows with the speed of change—.

Everyday substances do exactly this. Cornstarch mixed with water: strike it suddenly and it behaves like a solid —you can even run across it—; push it slowly and it flows between your fingers. Its resistance depends on how fast you force it. The Pleno has a similar character, and that "resistance that grows with speed" hides a deep consequence: it explains why nothing can go faster than light. Approaching the speed of light amounts to trying to deform the medium ever faster; and the medium hardens so much that it becomes impossible to push it further. The speed of light is not a rule imposed from outside: it is, simply, the point where the Pleno jams.

Why there is a speed limit

Every material has a curve saying how much it resists according to how fast it is deformed. The Pleno's is

$$n(x) = \frac{1}{1-x}, \qquad x = \frac{\text{rate of deformation}}{\text{maximum rate the medium admits}},$$

the curve of a dilatant fluid: resistance grows with x and, as x approaches one, it soars without bound.

There lies the content of the light limit. Accelerating a body means forcing the medium around it to reorganize ever faster, that is, raising x. The energy that costs goes as the resistance of the medium, and that resistance diverges at $x = 1$. It is not that pushing past c is forbidden by a rule: it is that it would cost infinite energy, which is exactly what the Lorentz factor says.

What reaching this point means. In special relativity, the speed limit is a postulate: it is stated at the outset and everything else hangs from it. Here nothing is postulated; the response of a material is computed and the limit appears where its curve turns vertical. One of the two cornerstones of twentieth-century physics goes from being a starting point to being a consequence. **[DERIVED]**/**[SLD]**

Two properties, then, two fruits: from *firmness* will come the weakness of gravity; from *resistance to abrupt change*, the limit of the speed of light. And from here come space and time as well. Space is the extension of the medium itself; and time is not a stage where things happen, but the very pulse of the Pleno. They are not the theater where the substance acts: they are aspects of the substance itself.

7. Time: not a stage, but the pulse of the Pleno

We are used to thinking of time as a river that runs on its own, independent of everything, and of the things of the universe as boats floating on it. First comes time; then, what happens within. In the Pleno that image is inverted, and the inversion is one of the deepest ideas of the whole proposal.

Here time does not exist "on the outside." There is no prior river. What we call time is, simply, how much the Pleno has changed: the count of its accumulated changes, measured against the only reference rhythm there is —the maximum speed at which the medium can react, the speed of light—. Think of a clock: a clock does not measure a time that would be there anyway; with its ticking, it makes time, counting it beat by beat. The Pleno is like that: it does not age within time; its aging is time.

This sounds abstract, but it dissolves at a stroke questions that seemed unfathomable. "What was there before the universe?" or "how fast does time pass?" turn out to be ill-posed questions —like asking what lies north of the North Pole—. There is no "before" without a Pleno that changes, because "before" is itself a measure of change; and there is no "speed of time," because there is no second external clock against which to time the first. Time is not a container that can be filled or emptied: it is the trace the substance leaves as it transforms.

And there is a feature of the mathematical development that crowns the idea. Three things that at first seem distinct turn out to be one: the curve along which the Pleno relaxes from its most tense state toward rest, the great clock that marks the history of the cosmos, and the very passage of time. They are not three synchronized clocks: they are the same object described three ways.

8. Gravity: not a force that pulls, but the medium relaxing

For centuries we thought of gravity as Newton described it: an invisible force one body exerts on another, pulling it at a distance. The Earth "pulls" on us; the Sun, on the Earth. It works wonderfully for calculation, but leaves an uncomfortable question: how does one body pull another that is far away, without touching it, across the void?

A century ago, Einstein overturned that image and answered, in large part, that question. In his general relativity, gravity is not a force that pulls: it is the *shape* of space. A mass deforms the space around it —curves it—, and other bodies, as they move, feel no pull at all: they simply follow the straightest path that curved space allows them. The usual image is a heavy ball on a taut sheet: it sinks the surface, and a marble passing nearby falls toward it not because the ball attracts it, but because the ground beneath the marble is tilted. Einstein thus dissolved the mystery of the "pull at a distance": there is no pull, there is geometry. The Earth does not tug at the Moon; it bends space, and the Moon follows the curve.

This is where PIU-Ψ comes in, not to contradict Einstein but to answer him the question he left open. Einstein tells us *that* space curves, and describes with exquisite precision *how* it does. But he does not say *what* space is, nor *why* it can curve: he takes it as the ultimate stage, given. PIU-Ψ takes the next step. Space is not the final floor of reality: it is a description of how the Pleno behaves seen from far away. What Einstein calls "the curvature of space" is, seen up close, the tension of the Pleno. Put in one phrase: for Einstein, geometry is the *starting point*; for PIU-Ψ, it is the *finish line* —something that emerges from the medium, not something to be assumed—. Einstein geometrized gravity; PIU-Ψ explains what that geometry is made of. And when one does the math, Einstein's equations reappear intact inside PIU-Ψ —they are not toppled, they are *derived*—.

In the Pleno, gravity *is not a force that pulls.* It is something deeper, and it arises directly from the tension we already know. A mass —a star, a planet, yourself— is, seen up close, a disturbance of the Pleno's tension: a region where the medium is altered from rest, and that alteration fades gradually with distance. Now put two masses close together: their two alterations overlap in the medium they share. And it "costs" the Pleno energy to sustain two separate tense zones; it is more economical when they come together and merge. By its own elasticity, the medium tends to relax that tension, and the way to relax it is to bring the two masses together. What we call "gravitational attraction" is not one mass pulling another: it is the medium between them relaxing its tension by uniting them.

Think of two drops of water on a table: when they get very close, they merge into one. It is not that one "pulls" the other; it is the tension of the liquid itself that, as their edges touch, unites them. The gravity of the Pleno is of that same family. And so Newton's old uncomfortable question does not even arise: nothing pulls across the void, because there is no void; there is a continuous medium that explains everything by contact, point to point.

A very firm medium ought to produce a very strong gravity. Just the opposite happens: it produces the weakest force in all of nature. The reason is clean once one understands what the "gravitational constant" really is —that number which fixes its intensity—. In PIU-Ψ that number does not measure a force: it measures **how much the medium yields** when you place a mass in it —and that yielding is what a distant observer reads as curvature of space—. And "yielding little" is exactly what being firm meant. Firmness and gravity are not, then, two facts: they are the same fact seen from its two sides. A medium that barely yields is a medium where a mass hardly alters the geometry read off it; and a minimal curvature is a minimal gravity. The gravitational constant is so small, precisely, because it measures the yielding of a medium that, at rest, barely yields.

Newton's constant as a structural identity

Take two modulus disturbances separated by a distance r and compute, with F1, the energy it costs to sustain them. The result has the form $-GM_1M_2/r$ of the Newtonian potential, and on comparing the two no loose factor remains: the coefficient multiplying F1's kinetic term *is* c^4/G. Solving,

$$\rho_P c^2 \ell_P^2 = \frac{c^4}{G} \qquad \Longleftrightarrow \qquad \boxed{G = \frac{c^2}{\rho_P \ell_P^2}}$$

That is: the gravitational constant is the speed of light squared divided by the density of the medium and by the square of its grain. Nothing else enters the formula.

This is not dimensional analysis. The program has dimensionless invariants of its own, and any of them multiplying that expression would give the same units; the derivation does not choose the form, it returns the **coefficient**, and the coefficient equals one. From it also comes the rigidity of the medium, $c^4/2G \approx 6 \times 10^{43}$ newtons, the same coefficient carried by the Einstein-Hilbert action.

One frontier remains. For G to give the exact number it gives in SI units requires having fixed the calibration of meter, kilogram, and second at one point. That is a statement about the measuring rod, not about the derivation —the SI value of a dimensional constant is underivable in principle, for any theory—. With that single point fixed, the structure forces the rest.

What reaching this point means. In general relativity and in the Standard Model, G enters as a datum: it is measured and accepted. No earlier theory returns it from its own dynamics. Here it comes out as an identity of the medium, and with it the smallness of gravity stops being an enigma: that smallness is the firmness of the Pleno read from the other side. **[DERIVED]**

And with it, Einstein's equations

The constant alone would not be enough. What is needed is for the whole equation of general relativity to reappear, and it does, by a known mechanism —Sakharov's—: integrate out the fluctuations of the medium's tensor mode, and what remains as the effective action of the geometry is, term by term,

$$\Gamma_{\text{ind}}[g] = -\frac{c^4}{16\pi G}\int d^4x\,\sqrt{-g}\,R + \mathcal{O}\big(R^2/E_P^2\big) \quad\Longrightarrow\quad \boxed{R_{\mu\nu} - \tfrac{1}{2}R\,g_{\mu\nu} = \frac{8\pi G}{c^4}\,T^{(\text{Pleno})}_{\mu\nu}}$$

The left-hand side is the curvature of spacetime; the right, what the Pleno contains at each point. It is Einstein's field equation, unretouched.

The coefficient comes out *complete*, not "of the order of": the factor $c^4/16\pi G$ appears with its value, with G already derived in the previous box. And from the same computation, solved around a static mass, one reads off the Schwarzschild metric with its radius $r_s = 2GM/c^2$ —the geometry of black holes—.

What reaching this point means. The curvature of spacetime is not a starting point here but the effective elastic action of a medium, and Einstein's equations are what that medium does seen from afar. General relativity stands intact inside the program, as an effective limit and with its exact coefficient: it is not corrected, its makeup is explained. **[DERIVED]**

None of this demands an act of faith: there is no need to assert beforehand "the Pleno is extremely firm." We can go the other way around. The weakness of gravity is a *measured* fact, one we all know; and from that measured weakness one *deduces* how firm the medium is, without having to assume it. The firmness of the Pleno is not a postulate to be believed: it is the fingerprint that gravity itself, weak as it is, leaves exposed. The weakness of gravity stops being a mystery and becomes the direct reading of that firmness at rest.

And how weak is it? More than intuition suspects. A simple fridge magnet, the size of a coin, lifts a paperclip overcoming the gravitational pull of the entire planet Earth pulling the other way. That pinch of magnetism beats the whole planet. And yet gravity, being trillions of trillions of times weaker than the other forces, is the one that governs the universe at large scales. How can the weakest rule over the galaxies? The key is its reach. The other forces, for all their intensity, either act only at tiny distances —the nuclear ones do not reach beyond the nucleus of an atom—, or cancel themselves out: in matter there are as many positive as negative charges, and their electric pulls almost entirely nullify each other. Gravity, by contrast, never cancels —there is no "negative mass" to counter it— and its reach has no limit: it adds and adds across cosmic distances, from planet to star, from star to galaxy. Thus the weakest force of all, precisely by acting without canceling over immense distances, ends up being the one that shapes the entire cosmos.

That colossal weakness of gravity is one of the great unsolved puzzles of physics: no one knows, in the usual framework, why it is so extraordinarily weaker than the rest. It is accepted as a fact, an "it is so." The beauty of a good theory is that it turns an "it is so" into an "it is so because"; and here the program replaces nothing —Newton's and Einstein's gravity remains intact—, but gives it an underlying reason: why its constant is as minuscule as it is. It is not the final word —much remains to confirm—, but it shows the aspiration: not only to describe the universe, but to understand why it is the way it is.

9. The other constant: the spin of a grain

If the firmness of the medium gives the gravitational constant, the other great constant of physics —the quantum one, Planck's— comes from the other side of the field: from the spin.

The Pleno has grain: it is not infinitely divisible, but made of cells of a minimum size, and each cell has an orientation of spin. A cell spinning on itself at the fastest rate the medium allows has a well-defined amount of rotation —what physics calls angular momentum—. That number, the angular momentum of a single grain spinning flat out, *is* Planck's constant. This is not an analogy: it is an identity, and it comes out of the same structure that gave gravity.

In the usual framework, gravity and the quantum are two separate theories, with two independent constants, and the problem of the century has been to marry them. In the Pleno, both constants are two readings of the same medium: one measures how much it yields when you squeeze it, the other how much its smallest grain spins. And there is a property of the medium —its stiffness against bending, something as elastic and as unquantum as the stiffness of a sheet— that turns out to equal exactly a quantum quantity. Gravity and quantum mechanics are not two theories to be reconciled: they are two ways of reading the same coefficient.

From the same structure comes Heisenberg's famous uncertainty principle, and it comes with its exact value: not "of the order of" Planck's constant, but with the factor one-half that quantum theory demands, and with a tiny correction that appears as one approaches the size of the grain.

Planck's constant is the spin of a grain

A cell of the medium has a moment of inertia about its internal phase, $I_\theta = \rho_P \ell_P^5$ —mass times radius squared, with the size of the grain— and can spin at most at the rate $\omega_P = c/\ell_P$, because nothing in the Pleno changes faster than c. Its angular momentum is the product of the two:

$$\boxed{\hbar = I_\theta\, \omega_P = (\rho_P \ell_P^5) \cdot \frac{c}{\ell_P} = \rho_P\, c\, \ell_P^4}$$

Again the coefficient the derivation returns equals one, and with a difference from the case of G: the density ρ_P was already fixed by the gravitational identity. Evaluating $\hbar$ with it is not a second fit, it is a consequence —and it reproduces the measured value to fifteen figures—. Between the two identities the relation $\ell_P = \sqrt{\hbar G/c^3}$ also holds exactly.

The stiffness of the medium against bending —F1's third term, a purely elastic property— turns out to equal $\gamma_\kappa = \hbar c/40$: a quantum quantity. And the uncertainty principle comes out with its exact floor, with the factor one-half quantum theory demands, plus a correction that appears only on approaching the size of the grain:

$$\Delta x \cdot \Delta p = \frac{\hbar}{2}\sqrt{1 + \frac{3}{80}\left(\frac{\ell_P}{\sigma}\right)^2}$$

What reaching this point means. Gravity and quantum mechanics have spent a century treated as two theories with two independent constants, and the central problem of fundamental physics has been to marry them. Here both constants come out of the same medium with the same two scales: one measures how much it yields when squeezed, the other how much its grain spins. There are not two theories to reconcile; there is one material read from two sides. **[DERIVED]**

Nor does it stop at the constant. Quantum mechanics is a theory with several statements accepted without proof: that the probability of finding a particle somewhere is the square of the wave function, or that correlations between separated particles may exceed what classical physics allows but stop at exactly a certain number. They are postulates: rules that work and that the theory does not explain. In the Pleno they cease to be.

The postulates of quantum mechanics, as consequences

Why probability is the square. The phase symmetry gives the medium an exactly conserved current —Noether's theorem again— and that current can be apportioned by region as a flow is apportioned. The probability of finding the particle inside a region is then not an added rule but the fraction of that current the region contains:

$$P(R) = \int_R |\psi_{\text{norm}}|^2 \, d^3x.$$

What in quantum mechanics is the Born rule is here the bookkeeping of a conserved quantity.

Why correlations stop where they stop. Two particles that shared an origin keep a phase relation, and measuring them at two angles a and b gives the correlation $E(a,b) = -\cos\big(2(a-b)\big)$. That factor of two —the phase turns twice as fast as the apparatus angle— is geometry of the medium's phase, and it drags an immediate consequence:

$$|S| \leq 2\sqrt{2} = 2.8284.$$

Classical physics cannot exceed 2; experiment reaches 2.83 and not a step further. That ceiling —the Tsirelson bound— is a measured fact quantum mechanics records without explaining; here it follows from how the phase turns.

And the equations of the quantum regime come out as limits of F1: at low energy the nonlinear Schrödinger equation of a condensate appears, its coefficients fixed by the medium; and on the diamond lattice the band structure produces cones with a complete Clifford algebra, whence Dirac's equation for the electron emerges by two converging routes, with no spinors postulated.

What reaching this point means. This whole branch —from $\hbar$ to the floor of uncertainty, from the Born rule to the bound on correlations— *contains not a single calibrated link.* There is no parameter in it where a discrepancy could be absorbed, which makes it the most exposed part of the program: any of these numbers could have come out wrong with nowhere to hide it. And from the same branch an astrophysical result is reached, the mass at which a star collapses,

$$M_{\text{Ch}} = \frac{\sqrt{3\pi}}{2}\big(-\xi^2\theta'\big)_{n=3}\frac{(\hbar c/G)^{3/2}}{(\mu_e m_u)^2} \approx 1.456\, M_\odot,$$

with $\hbar$ and G already derived, against the known value of about 1.44 solar masses. **[DERIVED]**

10. Where "everything" comes from

With the few ideas we already have —an elastic crystalline medium, with tension and spin, firm yet able to flow— the proposal reconstructs the entire catalogue of all that exists:

- **Particles** (a proton, an electron) are stable knots of the medium: twists that coil around themselves and do not easily undo, like a well-made knot in a rope. That is why they are stable, and why they come in whole amounts: a knot is either there or not. Their mass is, roughly, the energy stored in holding that knot.

- **Light** and all electromagnetic waves are waves of the spin running through the medium: the local orientation rotates from point to point, and that rotation travels. No separate electromagnetic field is needed; it is the Pleno itself undulating in its orientation.

- **Dark matter** —that invisible mass holding galaxies together— would be another family of configuration of the same medium. Ordinary particles are knots that, besides tensing the Pleno, make it spin, and from that spin electric charge and light are born. Dark matter would be a smoother lump: it weighs (it tenses the medium and so gravitates), but it lacks that spin —the circulation that produces electromagnetism—. And there is the reason it is dark: it is not that it "barely interacts" by chance, it is that it does not have the mechanism that generates light, so it cannot shine or reflect anything. It is not an exotic particle of another nature: it is the same Pleno, organized without that spin. Ordinary and dark matter are, then, first cousins —one knotted and luminous, the other smooth and mute—. The theory, moreover, *calculates* how much dark matter there must be relative to ordinary matter, and the number comes from the pure geometry of the crystal, falling within the measurements' margin of error, without adjusting anything.

- **Dark energy** would be the state of deep rest of the Pleno itself: even where there are no particles or waves, the relaxed medium keeps a basal energy, present everywhere, like the floor of the ocean over which everything else happens. In its state of rest —the state the universe is *already* in today— that energy behaves **exactly like a constant**: the same everywhere and at every moment, without the slightest variation. It is not "almost" constant: it is a genuine constant, because it is a symmetry property of the medium's rest, not a field still rolling somewhere. What does change over time is its *relative weight* in the balance of the universe, and the reason lies in what expansion is. What expands is not a container holding the Pleno: it is the Pleno itself distending. As it distends, the knotted configurations —matter— end up ever farther apart and their density falls; the resting floor, by contrast, keeps its own, because it depends not on how far the medium has distended but on where the bottom of its energy curve lies. The fraction the floor represents thus grows little by little, up to a ceiling fixed by the geometry of the crystal. That is why we see it dominate today: not because it increases, but because everything else thins out.

- And at the largest scale of all, **the entire Pleno could breathe**: compress and distend again in cycles. Since the medium cannot be compressed to infinity —recall its ceiling—, a collapse does not end in a dead-end catastrophe, but bounces and gives rise to a new expansion. And that bounce is not a distant detail: in the proposal it is the moment when everything else is born. As it decompresses and cools, the Pleno "sets" —like water turning to ice, forming crystals—, and in that setting the different configurations separate out and become fixed: the matter we see, dark matter, the dark-energy background, and even the few kinds of fundamental particles that exist. The bounce is, so to speak, the forge where the inventory of the universe was cast.

Light and the forces: the phase that circulates

Take the phase of the Pleno and travel once around an elementary face of the lattice, returning to the starting point. What in general does not return to its initial value is the angle: a net circulation remains. That circulation, written as a field, is electromagnetism:

$$E_{\text{mag}} = \frac{8}{9} F_{\mu\nu} F^{\mu\nu}, \qquad F_{\mu\nu} = \partial_\mu a_\nu - \partial_\nu a_\mu,$$

where a_μ is the phase angle accumulated in each direction and the coefficient $8/9$ comes from the geometry of those faces, not from a fit. From it follow the Bianchi identities and, via Euler-Lagrange, Maxwell's equations. If the internal group does not commute, the same construction adds the term $i[A_\mu, A_\nu]$ and Yang-Mills appears, the structure of the nuclear forces.

And light travels at c without having to impose it: the phase disturbance propagates at the maximum rate of change of the medium. The speed of light and the limiting speed are not two constants that happen to coincide; they are the same quantity measured two ways.

What reaching this point means. Maxwell and Yang-Mills are obtained without postulating any electromagnetic field: there is no separate field living on space, there is a single substance whose phase circulates. That the two gauge theories of the Standard Model come out of the same geometric construction, with a derived coefficient, is what allows treating them as branches of the same trunk rather than independent chapters. **[DERIVED]**

Matter: knots, whole charge, and three quarks

The same field admits configurations that cannot be undone by stretching: knots. They come in two classes according to how they wind —skyrmions, with a winding number B conserved by topology, and *Q-balls*, held together by the conserved charge of the phase— and that classification is the classification of matter.

Electric charge is the winding number. Not a property assigned to the particle, but how many times the phase wraps around the knot. And a winding can only be a whole number: there is no half turn. From this follow two exactitudes that standard physics takes as given without explaining them. The neutrality of matter —that the charges of proton and electron cancel— is exact by construction, and experiment confirms it to one part in 10^{21}. And the proportionality between charge and mass holds to one part in 10^{15}.

Chemistry comes from here too. When the winding number is odd, a topological theorem —the Finkelstein-Rubinstein criterion— forces the knot to behave as a fermion. That is: the Pauli exclusion principle, and with it the fact that electrons stack in shells and a periodic table exists, is not a postulate added to quantum mechanics but a consequence of how the medium is knotted.

And the three quarks. The proton's $B = 1$ winding is not shared out in any old way: it splits into three Borromean-linked vortices —three rings that hold one another and come apart if any one is removed— and arithmetic modulo 3 over those three returns the observed charges: $(2, 2, -1)$ for the proton, $(2, -1, -1)$ for the neutron. With the proton mass as the only input scale, the size comes out $R_p = 0.851$ fm against 0.8409 measured, and the Δ mass 1202 MeV against 1232.

What reaching this point means. That matter comes in whole units, that charge is quantized, that quarks are three and not two or five, and that an exclusion principle exists, are facts the Standard Model describes precisely but does not explain: it takes them as given structure. Here they are consequences of the medium being knotted, and that the computation also closes on the proton radius to within one percent indicates the picture is not merely qualitative. **[DERIVED]/[SLD]**

All of this, with *a single substance*. There is no ingredient for gravity, another for light, another for dark matter. There is the Pleno, and its different ways of tensing, spinning, undulating, and flowing.

The share-out of the cosmos, from the geometry of the crystal

Ordinary matter knots itself and therefore occupies the whole sphere of directions of the medium; dark matter, lacking that knotting, occupies only the cell of the lattice —its Voronoi tetrahedron—. The proportion between them is then a ratio of two areas, and that ratio is a pure number:

$$\boxed{\frac{\Omega_{DM}}{\Omega_B} = \frac{2A_{\text{sphere}}}{A_{\text{tetrahedron}}} = \frac{2\,(4\pi R^2)}{(8\sqrt{3}/3)R^2} = \sqrt{3}\,\pi \approx 5.4414}$$

The Planck 2018 measurement gives 5.375 ± 0.077. That is agreement at $\mathbf{0.86}\sigma$ with *zero* free parameters. And the comparison is clean: it is made against physical densities read directly off the heights of the cosmic background peaks —where the h^2 factor cancels— not against parameters that are already the output of a previous fit.

The individual values of each component further inherit a state quantity —where the medium stands today along its path—, which in this sector reads as the condensate fraction f_c. It is the same quantity that sets the hour on the cosmic clock: one state reading across the whole program, not two calibrations. With it, the whole block of densities falls out:

Quantity	PIU-Ψ	Planck 2018	Agreement
$\Omega_{DE} = f_c$	0.6869	0.6847 ± 0.0073	0.30σ
$\Omega_m = 1 - f_c$	0.3131	0.315 ± 0.007	0.27σ
Ω_{DM} (dark)	0.2645	0.2645 ± 0.007	0.07σ
Ω_B (baryonic)	0.0486	0.0493 ± 0.0006	1.2σ

Dark energy, moreover, is the phase condensate at rest at the bottom of the well, whence its energy tensor is proportional to the metric: $w = -1$ *exactly*, with no residue. And any residue could have only one sign —positivity of the kinetic term imposes $w+1 \geq 0$— so the theory vetoes the phantom region where w would be below -1.

What reaching this point means. The dark sector is 95% of the content of the universe, and in the standard model its proportions are parameters fitted to observation: nobody knows why they take the values they take. Here the ratio of dark to ordinary matter has no knob at all: it is a quotient between the area of a sphere and that of a tetrahedron, and it lands inside the satellites' error bars. The absolute values hang from a single state quantity, the same one throughout the program. A number taken from the geometry of a crystal getting the composition of the cosmos right is the kind of coincidence that does not happen by chance. **[DERIVED]**/**[CALIBRATED]**

That leaves the last item on the list, the bounce, which is where everything above was forged.

The bounce: why there is no singularity

Compressing the Pleno means climbing the upper wall of its well, and that wall diverges. The consequence can be quantified: the self-gravity of a collapsing region falls short by a factor of some four hundred thousand of the energy needed to reach the wall. The saturation configuration is *unreachable*, and not narrowly so.

What happens instead is a bounce, and the equation governing it is that of a damped oscillator:

$$\rho_P \ell_P^2 \, \ddot{\bar{S}} + \eta(\bar{S}) \, \dot{\bar{S}} + V'_{\text{eff}}(\bar{S}) = 0,$$

where $\bar{S}$ is the mean modulus of the universe, V'_{eff} the slope of the well —the force returning it to the bottom— and η the damping of the medium. That damping is not constant: it grows on compression, from 0.984 in the present configuration to some 10^4 times more against the wall. An overdamped bounce is a clean bounce, with no resonance or ringing.

What reaching this point means. The initial singularity is, in general relativity, a theorem: under very general conditions collapse ends at a point where the equations cease to make sense. Models that avoid it normally do so by adding exotic matter or choosing particular initial conditions. Here it is not avoided: it is *inaccessible*, because the material does not admit that compression, and the bound preventing it is the same A6 wall that gave the ceiling of the well. The beginning of the universe stops being a point where physics breaks and becomes the point where the medium turns around. **[DERIVED]**

11. The heat of the medium itself

One whole family of laws remains to be explained, and it is the one that asks about heat. What is temperature, if all that exists is a single substance? In ordinary physics, heat is handled with its own apparatus —thermodynamics and statistical mechanics—, with its own rules about how many particles share out energy. Here that loan is not available: if the Pleno is all there is, heat has to be something that happens to the Pleno, and its laws have to come out of counting it.

And they do. Temperature, on this reading, is disorder of the spin: when the orientations of all the cells are in step, the medium is cold —that is the condensate, the rest state, dark energy—; when each cell spins on its own, out of agreement with its neighbors, the medium is hot. Temperature and decoherence are the same thing. And since the medium has a compression ceiling, it also has a maximum temperature: the point where the gas of vibrations reaches the density of the medium itself and no more heat fits.

Without assuming any statistical distribution in advance —without importing anything from the physics of heat— it is enough to share out indistinguishable vibrations among the modes of the medium and demand equilibrium for Bose-Einstein statistics, which governs photons and the whole family of bosons, to come out on its own. And integrating that occupation over the modes gives the Stefan-Boltzmann law, the one that says how much a hot body radiates, with its exact coefficient. The thermodynamics of the Pleno is not borrowed from another theory: it is deduced from counting its own configurations.

Bose-Einstein and Stefan-Boltzmann, by counting

Three steps, and in none of them is a statistical distribution assumed. **First**, expand F1 to second order about rest: each phase mode becomes a harmonic oscillator, its frequency ω_k fixed by the medium. **Second**, count in how many ways *indistinguishable* vibrations can be shared among those oscillators; it is pure combinatorics, and gives the entropy per mode $\sigma(n)/k_B = (1+n)\ln(1+n) - n\ln n$. **Third**, demand equilibrium using the native definition of temperature, $1/T = dS/dE$. That forces the mean occupation:

$$\boxed{\langle n\rangle = \frac{1}{e^x - 1}, \qquad x = \frac{\hbar\omega}{k_B T}}$$

which is Bose-Einstein statistics. And the bosonic character is not postulated either: the phase mode is a Goldstone boson, its momentum comes in whole multiples of $\hbar$, and it has no occupation ceiling.

Integrating that occupation over all modes gives the entropy per unit volume, which is what fixes how much a hot body radiates:

$$\frac{s}{k_B} = \frac{g_*}{2\pi^2}\int_0^\infty k^2\sigma_k\,dk = 0.877298168986 \qquad \text{against} \qquad \frac{2\pi^2}{45}\,g_* = 0.877298168986,$$

the coefficient of the Stefan-Boltzmann law, agreeing to thirteen figures.

What reaching this point means. Thermodynamics and statistical mechanics are, in ordinary physics, a separate apparatus with its own postulates about how energy is shared. Here none is imported: the statistics of photons and the law governing the brightness of a star come from counting configurations of the same substance that gave gravity and charge. A whole branch of physics believed to be independent turns out to be a consequence of the same Lagrangian. **[DERIVED]**

The same computation yields a dissipation floor —a minimum viscosity below which the medium cannot flow— and it has a clean physical reading. In a dilute gas, nothing prevents a fluid from flowing with no internal friction at all: it is enough to dilute it. In the Pleno it is impossible, because the medium is always full and coupled —there is nowhere to dilute, there is no vacuum—. The existence of a minimum viscosity, conjectured in ordinary physics from string theory, is here the impossibility of vacuum read as a transport property. And the window the theory predicts is the band where the quark-gluon plasma and ultracold atomic gases actually sit in the laboratory: the same medium in regimes separated by twenty-one orders of magnitude in temperature.

12. How the Pleno closes the gulf between the large and the small

The great problem we began with was this: the very large (gravity, the cosmos) and the very small (particles, the quantum) are described by two theories that do not fit, and no one has managed to unite them in a century.

PIU-Ψ does not try to glue the two theories together: it starts from the premise that they were never two things. If everything —the large and the small— is the same Pleno, then gravity and the quantum are not two worlds to reconcile by force, but two behaviors of the same medium at different scales. At the small scale, the Pleno shows up as particles and quantum phenomena; at the large scale, as gravity and cosmology. There is no gulf to bridge, because beneath both lies the same substance.

And this explains, in passing, why PIU-Ψ sidesteps the jam of infinities that ruined all earlier attempts. That jam arises from treating gravity as a fundamental force and trying to "quantize" it. But in the Pleno gravity is not fundamental: it emerges from the medium at large scales, as we saw. There is no need to quantize it, just as we do not "quantize" the waves of the sea —we understand they are water—. And since the medium has a ceiling (it does not compress to infinity: it has a ceiling), the infinities simply do not appear: the material itself sets the limit. That a single idea speaks naturally of both the cosmos and the atom does not solve the old problem at a stroke —work remains to be done—, but it does something more elegant: it dissolves it, showing that perhaps there were never two things to unite, but one seen from two distances.

The two modes of the medium, and where the cutoff comes from

Disturbing the Pleno slightly gives two waves that do not mix: one in the modulus —how tensed it is— and one in the phase —which way it spins—. That they are decoupled at first order is the structural reason standard cosmology can treat the universe as a sum of independent fluids and have it work.

The modulus mode is *massive*: its frequency never drops below a minimum,

$$\omega^2 = c^2k^2 + \omega_0^2, \qquad \omega_0^2 = \frac{\xi\, c^2}{\ell_P^2}, \quad \xi = 19.69,$$

where ξ is the curvature of the well at its bottom, computed from the potential. That mass is of the order of the grain, enormously high, and the mode has been frozen since crystallization: no background dynamics remains that could move w from -1.

The phase mode travels at the speed of light, with a tiny correction noticeable only at wavelengths near the grain:

$$\omega_{\mathrm{TT}}^2 = c^2k^2\left[1 + \frac{(k\ell_P)^2}{20} + \cdots\right],$$

the same correction that affects the photon. And in the ultraviolet, the resistance to being bent —the flexural term— makes the second branch of the propagator an overdamped pole, of the type studied by Lee and Wick. There is the cutoff: the medium admits no curvature tighter than its grain, and that refusal is what keeps the computations from diverging.

What reaching this point means. The infinities that wreck attempts to quantize gravity are normally headed off with a cutoff introduced by hand and justified afterwards. Here there is no cutoff to introduce: the material has grain, and that grain was already in F1 from the first axiom. Ultraviolet regularity is not a remedy of the formalism, it is a property of the substance. **[DERIVED]**

13. Why this is science, and not a pretty tale

A beautiful story is not enough. What makes this a serious proposal and not an agreeable metaphor? That beneath every image in this document there is a formal research program, governed by the same demands as any professional physics. The Pleno is not described with words, but with **high-level mathematics**: it starts from precisely stated principles, derives its results the way one proves theorems, subjects every step to verification —recomputing the numbers with an independent program that carries no preloaded value— and classifies every claim by its degree of certainty, without mixing the proven with the conjectured. And with that apparatus, the theory does three things a mere speculation never could.

It reproduces laws we already know to be true. That long dozen of laws from the start are not guesses: the theory derives them —it does not adjust them, it derives them— from its single equation. When PIU-Ψ reaches one of those laws, it is not proposing something new to check, but reaching by another route a result already confirmed. And a validated law does not lose its validity for being reached from another starting point. This part of PIU-Ψ, therefore, rests on old, solid science, observed in thousands of experiments. **Its contribution is not to replace that science, but to integrate it**: to give laws that today live in separate books a single origin from which all of them follow.

It hits where there is no room to accommodate. The program's successes are of two kinds, and both are hard. The first are *exact coefficients*: the theory does not reach a law "up to a factor," but with the full number in front. The floor of uncertainty comes out with its factor one-half; the Chandrasekhar limit, with its entire coefficient; the bound on quantum correlations, exact; the time-stretching of distant supernovae, exact and confirmed by the data. The second are *predictions without parameters*: the proportions of the cosmos —how much dark matter, how much ordinary, how much dark energy— come out of the geometry of the crystal and fall within one standard deviation of what the satellites measure, with not a single knob to turn.

In the entire tree of results there is **exactly one calibrated ingredient**, and it is declared: the present position of the medium along its path —the analogue of the hour a clock reads—. It shows two faces, the condensate fraction of the dark sector and the Hubble constant, but it is one quantity read in two projections, not two fits. From it hangs the whole cosmological block, with no further adjustments. No coefficient of the theory is fitted to the data; no quantum branch contains any calibration; and the single anchored quantity is a state datum, not a free parameter.

And it hits not once, but many times —and that changes everything. Reaching a known law from a new origin could be lucky chance. But PIU-Ψ reaches more than a dozen, toward phenomena with nothing to do with one another. Here the argument takes a leap: a key that opens one lock might have forced it by luck; but the same key that opens twenty different locks is no longer chance: the key is the right one. When a single idea reproduces the quantum world, the mass of the stars, electromagnetism, the radiation of a hot body, and the proportions of the cosmos, at some point it stops being reasonable to speak of coincidence. Scientists call this *unifying power*: the more unrelated phenomena fall under the same explanation, the harder it is to believe it is an accident.

That argument is here put into numbers. The natural suspicion about a theory that hits with short formulas is numerology: with enough combinations of π, roots, and small integers, something will eventually coincide with something. That suspicion can be turned into an answerable question —how many simple formulas would land on these values by chance?— and answered by enumerating the entire space of formulas with a computer. The more ambitious question can also be asked: how many complete theoretical frameworks would reproduce the *whole* set of facts, and not just one lucky ratio? Both questions have quantitative answers, and both point the same way.

The two Bayes factors, and what each one says

They are two distinct questions, which is why their factors are neither multiplied nor averaged.

The narrow objection: how many *formulas* would land there? The entire space of algebraic expressions constructible from $\{1, 2, 3, 4, \pi, e\}$ up to five operations is enumerated by computer —3778 distinct expressions— and one counts how many fall inside the 2σ window of the observed values. The probability that two of them hit both $\sqrt{3}\pi$ and f_c at once is below 3×10^{-4}: a factor of some 6.7×10^5 against reading them as lucky formulas.

The broad objection: how many *frameworks* would reproduce the set? Fifteen results are admitted under four criteria declared *before* counting —that the outcome discriminate, that there be no adjustment freedom, that the target was established beforehand, and that the measurement channel be its own—, read with fifteen instruments that share no mechanism. Another twelve results are left *out* of the count, with their reason published: not for being incorrect, but for not contributing independent evidence. With the product of probabilities at 8.4×10^{-10} and a penalty for the shared trunk:

$$\mathrm{BF} \approx 3.4 \times 10^4, \qquad \text{range } [\,65 \text{ to } 1.2 \times 10^9\,]$$

depending on how severe that penalty is made. It is always reported with the range, never as a point.

Robustness. Tripling *all* the probabilities and applying a harsh penalty as well, the factor stays above 10; under the most punitive assumption —equivalent to declaring the fifteen results almost a single fact— it still equals 65; and removing one item at a time, no individual omission drops it by more than a factor of four and a half.

What reaching this point means. A Bayes factor compares two hypotheses and validates neither:

this does not prove the program correct. What it does establish is that the alternative explanation —that the successes are numerological coincidences— is disfavored under any reasonable choice of assumptions, including ones designed to hurt the program. The most natural suspicion about a theory that hits with short formulas is thereby answered with a computation, not with a retort. [METHOD]

To this is added, as an argument of another nature and therefore not summed into the factor, the coherence across sectors: that the same medium, with the same two scales, yields at once the gravitational constant, the quantum constant, the exact Heisenberg floor, and the mass at which a star dies. And it all rests on a telling fact: many of those numbers spring, inevitable, from the simple tetrahedral geometry of the crystal —the angle of the bonds, the number of neighbors—, not from a convenient adjustment.

Finally, in PIU-Ψ there is not a single freehand claim. Each result is developed with its explicit mathematics and carries, stated exactly, its degree of certainty. That discipline is what gives content to the word *derived*. Where the program has the structure but not the closed value, it says so; where it uses a result of others, it marks it as an import at the point where it enters, and does not recalibrate it.

How each claim of the program is to be read

Every result carries one of these labels, and none is used for another:

[DERIVED]	complete proof, with no conceptual gap
[SLD]	strong logical deduction: the structure is there, the closed algebra is not
[CALIBRATED]	value anchored to a datum, with the anchor declared

Numerical verification issues **two verdicts that must not be confused**. The first is document-code fidelity: that the program recompute exactly what the document claims, with zero tolerance, and it must pass in full. The second is agreement with measurement, which applies only where there is something to measure —a structural identity admits no agreement figure, since publishing one would present algebra of the input as if it were prediction—.

The program, in numbers

From a single statement —the universe is finite, conserved energy— and a single dynamical object, F1, PIU-Ψ derives **72 results organized into eight chains**, joined by **69 links** that follow one from another without inserting new assumptions along the way. Of these, **41 quantities are recomputed numerically** from their first-principles formulas, with no preloaded values; the verification notebook recomputes **99 quantities** in total —those forty-one plus their intermediates and controls— and all of them pass.

In that whole tree there is **one single calibrated ingredient**, a state quantity, declared; **three imported results** from the literature, marked as such and not presented as native derivations; **seven falsifiable predictions**; and **fourteen open loose ends** inventoried one by one with their severity.

That verification is not asked for on faith. The computer program that recomputes each of those quantities is available for download on the project's website, so anyone can run it and check the numbers for themselves. And behind every claim there is a complete formal development: a corpus of more than four hundred pages where each derivation is done step by step, and a mathematical compendium that condenses the proofs with their intermediate algebra and unit analysis. The invitation, as always in this program, is to verify.

14. The language in which the universe operates

That last fact —that the numbers are not arbitrary but spring from geometry— has a deep consequence. Taken together, all the numbers the theory produces form a single pattern:

geometric ratios born of the crystal's shape, combined by fixed rules. Put another way, the apparatus of the program contains a kind of "language" with which the Pleno generates its magnitudes: a grammar, with its elementary words and combination rules, from which the values of the theory follow.

In ordinary physics, the constants of nature are measured and accepted as they are, without knowing why they are worth what they are. The language of the Pleno changes that relationship: it lets one read why a number is what it is, and even anticipate the value of a magnitude from its place in the grammar, before computing it another way. A grammar does not only describe sentences already written: it lets one write new ones. It is the difference between memorizing a list of constants and understanding the rule that produces them —between reading the universe and understanding the language in which it is written—.

What a pure number is, and why it matters

Almost every quantity in physics carries units and, with them, a measuring rod chosen by convention: a meter, a kilogram. A *pure number* is a quantity the medium produces with no rod at all —any civilization, in any system of units, would obtain the same number—. That the program generates them is what allows speaking of a grammar rather than a list of measured constants.

The clearest example is the dissipation floor of the medium. The computation runs through the maximum temperature, Planck's constant, and Boltzmann's, and on substituting $\hbar = \rho_P c \ell_P^4$ all the constants cancel against one another —k_B included, because in the native language of the medium temperature *is* energy— and what remains is

$$R_{\text{floor}} = \frac{2 \cdot 5^{3/4}\, 6^{7/8} \sqrt{\pi}}{75} = 0.7579726,$$

with not a single free symbol. The same happens with the depth of the well, $E_0 = 4.2546$, closed in π and $\sqrt{6}$, and with the point at which dark energy equals matter, $\sqrt{2}\, S_{\min} = 1.0226$ —the geometry fixes in which *state* that crossing occurs, though not when—.

What reaching this point means. When a constant comes out as a pure number, it stops being a datum that must be measured and becomes something that can be *read*: one sees which geometry it comes from and why it could not have been otherwise. That is the difference between having a table of constants and having the rule that produces them. **[DERIVED]**

15. The cosmic clock: reading the hour of the universe

There is a number astronomy measures with care: the speed at which the universe expands, known as the Hubble constant. In ordinary physics it is, once again, a datum that is measured and accepted —one more number in the list of "it is so"—. The Pleno's proposal offers a different and beautiful way to understand what that number is.

The key is an idea we already saw: in the Pleno, time is the pulse of the medium, counted in tiny "ticks" —the smallest physics knows, the so-called Planck ticks—. The theory shows that the expansion speed of the universe is not a fundamental ingredient, but something the theory *computes*. PIU-Ψ gives a closed formula for that speed, with three pieces: the geometry of the Pleno, the rhythm of the Planck tick, and the state the medium is in today. The first two come entirely from the theory; the third is the reading of the cosmic clock at this moment. And when that formula is evaluated, it gives the number: about **67 kilometers per second per megaparsec** —exactly the value the satellites measure—. It is not an adjustment nor a datum put in by hand: it is the result of a calculation, and the calculation is right.

One nuance is worth stating, because it is subtle and beautiful. That number is not an eternal constant engraved in nature, but the reading of a clock at a given instant —the "hour" of the cosmos—. Asking "why is the expansion exactly that and not something else?" is like asking a clock "why is it 3 and not 5?": the clock's mechanism is the theory's business, and the theory has it; the particular hour depends on where the universe is along its path. The whole curve is theory; the point on the curve, today, is what the clock reads now. This —subtle as it seems— defuses one of the greatest scandals of modern physics. There is a number famous for its absurdity: the vacuum energy that quantum theory predicts exceeds the observed one by a factor of one followed by *a hundred and twenty zeros*. It has been called "the worst prediction in the history of physics." Under the Pleno's reading, that monster is tamed: it is not an impossible fine-tuning, but, once again, the age of the universe in ticks, squared. A gigantic number, yes, but for an understandable reason and not a miraculous coincidence.

The hour on the clock, and three anomalies reduced to one fact

The expansion rate is not an ingredient of the theory but something the theory computes, and its formula has three pieces:

$$H_0 = \sqrt{\frac{8\pi}{3}}\,\frac{c}{\ell_P}\,\sqrt{\Theta_0}\;,$$

a geometric constant, the rate of the Planck tick c/ℓ_P —the fastest clock there is— and Θ_0, which is where the medium stands today along its path —the same state quantity that appeared in the share-out of the cosmos as f_c, read here in another projection—. The first two are pure theory; the third is the reading of the clock at this moment, and it equals

$$\Theta_0 = \frac{3}{8\pi N_{\text{tick}}^2} \approx 1.66 \times 10^{-123}, \qquad N_{\text{tick}} \approx 8.5 \times 10^{60},$$

where N_{tick} is simply the age of the universe counted in Planck ticks. Evaluated, the formula gives the 67 km/s per megaparsec the satellites measure.

With the position fixed, the expansion history is forced by the potential, with no freedom to retouch it. Against the DESI measurements it gives $\chi^2 = 6.54$ against 6.53 for the standard model —statistically identical, with the same degrees of freedom, but with parameters derived rather than fitted—. And the time stretching of distant supernovae comes out exact, $b = 1$ against the observed 1.003 ± 0.011: agreement at 0.27σ with no parameters at all.

What reaching this point means. That 10^{-123} is the number physics calls the worst fine-tuning in history: the predicted vacuum energy exceeds the observed one by that factor. Here it is not a tuning, it is the inverse square of the age of the universe in ticks —a large quantity for an understandable reason—. And the same positional reading settles at a stroke two further anomalies the standard model treats separately: the primordial H_0 is not a measurement but the value that model needs to fit the cosmic background, and DESI's variable-dark-energy signal is what forcing a medium with $w = -1$ into a mold expecting the opposite produces. Three open problems, one single fact. **[DERIVED]/[SLD]**

And from here comes a strong claim about a puzzle that vexes astronomers. Measuring the expansion speed two ways —one looking at the young, distant universe, the other at the near, present one— yields different numbers: about 67 in the first case, about 73 in the second. They call this discrepancy "the Hubble tension," and in the usual model no one knows how to explain it. PIU-Ψ says two things. **First**: the true background value of the universe is the low one, those 67 the theory already computed and that match what the satellites measure; the high value, 73, is ruled out as a background value —not by preference, but by an impossibility of the medium, which relaxes toward its rest always in the same direction, without going backward or putting on bursts of speed, and a background that high would demand precisely what is forbidden—. To be clear and avoid confusion: the 67 is not in doubt; it is a real number, the one the Pleno computes and the one observation confirms. What PIU-Ψ calls into question is not that value, but *that the two measurements should be incompatible.*

Second, and here is the program's own twist: that incompatibility —the "tension"— is not a physical anomaly demanding new matter or forces, but an *artifact of the measurement method.* When a medium like the Pleno is analyzed with the standard model's template —a template that is not its own—, the two ways of measuring seem to clash even though the background is a single one: the gap between 67 and 73 appears because of the mold through which the data are read, not because the universe hides two competing rates of expansion. Put another way: what is real is the 67; the "tension" is the mirage. From that comes a falsifiable prediction —among the most valuable things a theory can offer—: that discrepancy must *dissolve* when the expansion is measured with methods that do not presuppose the standard model, leaving a single coherent value. If instead the gap persists in those clean measures, the Pleno's reading would be in trouble. The same holds for the "changing" dark-energy signal that surveys like DESI begin to hint at: the program holds that it is the same kind of mirage —the reflection of forcing a constant dark energy into a mold that expects a variable one— and that it will vanish in measures that do not presuppose the result.

Several pieces are *settled* —derived or strongly deduced— at once: that H_0 is the hour on a clock and not an eternal constant; the reformulation of that "one followed by a hundred and twenty zeros" as the age of the universe in ticks squared; the exclusion of 73 as the background; and that the time-stretching of distant supernovae comes out exactly as it should, confirmed by the data. All of that is what the standard model *cannot* do, and the Pleno can. The only thing that remains *open* is a fine, bounded piece: the exact form of a small correction to the expansion curve in its most extreme stretch —a loose end of calculation, not a crack in the edifice—. The reading of the clock, with that well-delimited reservation, is one of the most distinctive results of the whole program.

16. How far we have come, and how far we can go

PIU-Ψ is the work of an independent researcher, a program still under construction, not yet a consensus of the scientific community, and one that openly acknowledges its open pieces and its loose ends. But to stop there would be as unfair as to exaggerate: this is not a loose speculation, but a mathematical edifice that already reproduces confirmed laws and already gets it right, with very high precision and without adjustments, wherever it can be compared with reality.

The program draws its own frontiers, and draws them by name. **What is calibrated** is a single quantity: the present position of the medium along its path, the analogue of H_0 in any cosmology. **What is imported** is three results coming from the literature —a Klauder coefficient, a Lane-Emden constant, a Skyrme moment of inertia—, used and marked as imports, not presented as native derivations. **What is open** is fourteen loose ends declared with their severity, from the weightiest —the cascade that would fix the lepton masses— to the slightest. Knowing exactly where those three frontiers lie is what makes it possible to take everything inside them at face value.

What reproduces already-validated physics is waiting for nothing: those laws are already true, and that PIU-Ψ reaches them by a new route does not put them in doubt —that base is solid today—. What time and data will confirm or refute are its own predictions, and they are of the kind a model with the same degrees of freedom cannot make: the proportions of the cosmos; dark energy as a genuine constant, whose relative weight grows toward a ceiling fixed by the geometry; the reading of the Hubble tension and the DESI signal as method artifacts that must dissolve in clean measures; an ultraviolet dispersion shared by the photon and the gravitational wave, within reach of LISA and the third-generation detectors; and a definite density —with no singularity— for the core of a primordial black hole. And, above all, its integral vision: that all those scattered laws are, in reality, faces of a single substance. That is the great wager. If it proves right, we will have discovered that the entire universe is a single substance tensing and spinning, and that the laws we believed separate were always the same. If not, its new predictions will fail and we will know —but the laws it already reproduces will remain true, because their validity never depended on PIU-Ψ—. In any case, it poses one of those questions worth asking, one that invites us to look at the sky and at matter with different eyes: what if everything —truly everything— were the same?

All together: an idea worth pursuing

We began from an uncomfortable fact: physics, for all its triumphs, is still a puzzle with pieces that do not fit —the very large and the very small described by theories that contradict each other, and nearly 95 % of the universe made of something we know is there but do not understand—. These are real, acknowledged gaps that have resisted for decades. And against them, PIU-Ψ dares a different question: what if they were not many badly-fitting things, but a single thing badly understood?

That audacity, on its own, would prove nothing. What makes PIU-Ψ worth attention is how it pursues that question: not with a comfortable tale, but with mathematics that reproduce already-confirmed laws, with high precision and no hand-tuned numbers, and exposing itself to refutation by the observations to come. This is exactly how science has taken its greatest leaps: someone dares to suspect that separate phenomena are manifestations of something deeper, and then submits it to the implacable test of the data. Newton united the apple and the Moon; Maxwell united electricity and light; PIU-Ψ attempts to unite, no less, the quantum and the cosmic under a single substance —not repealing Newton, Maxwell, or Einstein, but showing that all of them were writing, without knowing it, chapters of the same book—.

It is no accident, then, the name the program bears: *Universal Integrity Principle*. Each word says something. It is a *Principle* because everything springs from a single root —an equation, a starting point— from which the rest is deduced, rather than a catalogue of loose rules. It speaks of *Integrity* in its double, most beautiful sense: that of unity —a whole universe, entire, not fragmented into pieces that do not talk to each other— and that of coherence —an edifice that does not contradict itself nor break into infinities—. And it is *Universal* because it reserves no corner: it aspires to embrace all that exists, from the tiniest particle to the entire cosmos, with a single explanation. The name, in short, is not a label: it is the very motivation of the program made word —the search for the whole unity of all there is—. And the symbol Ψ that appears beside its initials recalls where everything starts: the Pleno, written as the wavefunction of a universe that pulses, at bottom, like a single thing.

Where to go next

All the material of the program is available for download at **piuniversal.com** and on Zenodo, under a Creative Commons license.

PIU-Ψ V5.1 corpus	the complete program, each derivation step by step (425 pp.)
Math compendium	the intermediate algebra and the unit analysis (107 pp.)
Verification notebook	the program that recomputes the quantities, with its tables

The Spanish and English editions are equivalent: same structure, same equations, same classification result by result. A *Synthesis of results* —the inventory of the eight chains, result by result— accompanies the Spanish edition. **DOI: 10.5281/zenodo.21724390**

piuniversal.com • Manuel Alberto Celedón Mejía • Medellín, Colombia

Zenodo DOI: 10.5281/zenodo.21724390