Precausal Substrate Theory (PST)

On the Nature of Reality, as I see it

Ralf Meelker, 2026 - ralf@meelker.nl +316-11.122.133

Abstract

Precausal Substrate Theory (PST) proposes that reality does not originate within spacetime, but arises from logic itself. The theory posits a nonspatiotemporal, precausal substrate with a single irreducible primitive, **property differentiation** (the capacity for distinctions to exist at all), whose first expression is **asymmetric tension**. Once this tension exceeds a critical **modal threshold**, the configuration undergoes **modal sublimation**: a direct, nontemporal phase transition into **instantiated geometry**, from which spacetime and causality emerge together as coemergent consequences. Because loci are not primitive, PST requires no singular origin event; the Big Bang is a perspectival artifact of geometry that was never born at a point. The cosmological implication is a continuous, unbounded field of instantiation whose higher dimensions level off toward four as tension redistributes. The resulting vacuum is a ring of degenerate minima; motion along it is the only stable state it admits, making orbital mechanics at all scales a structural consequence of vacuum topology rather than an assumption. Quantum Mechanics [23, 24, 25] and General Relativity [1, 10] both emerge as projections of the same precausal structure. The Einstein field equations (the equations governing how matter curves spacetime) are not postulated in PST but derived: the way the substrate's order parameter varies across configuration space projects onto spacetime as gravitational curvature, and a theorem of differential geometry (Lovelock's theorem) establishes that in four dimensions at low energies this is the unique form gravity can take, given the diffeomorphism invariance that PST derives from its own structure via Noether's second theorem. The equivalence principle, the empirical fact that all objects fall identically in a gravitational field, regardless of their composition, is not an assumption but a theorem of PST: because all matter and all curvature are projections of the same single source through the same operator, there is no structural room for them to respond differently to each other. A concrete empirical prediction follows from the theory's dimensional reduction: the substrate is not infinitely smooth but has a finite grain size $d_0 \approx 7\,\text{nm}$, and this granularity modifies the Casimir effect, the tiny attractive force between two uncharged metal plates held very close together, producing a correction that scales as d^{-6} rather than the standard d^{-4}. This power-law exponent is a categorical, parameter-free prediction fixed by the symmetry of the substrate alone; at plate separations of 50 nm the correction reaches $\sim$2%, within reach of near-future precision experiments. By grounding the emergence of spacetime in a logically necessary modal threshold, PST offers an answer to why there is something rather than nothing.

v9.9

Contents

v9.9

v9.9

v9.9

1. Introduction

1.1. The incompatibility of quantum mechanics and general relativity

The search for a unified account of physical reality has long been constrained by the assumption that spacetime, matter, and causal order constitute the fundamental architecture of the universe. Modern physics inherits this assumption from classical ontology, even as its two most successful frameworks, Quantum Mechanics and General Relativity, demonstrate that this assumption cannot be sustained. Quantum Mechanics describes a domain in which discreteness, nonlocality, and the breakdown of temporal sequence are intrinsic features of physical behaviour. General Relativity, by contrast, models the universe as a smooth, continuous geometric manifold whose curvature governs gravitational interaction.

These frameworks do not merely differ in mathematical form; they presuppose incompatible ontological primitives. General Relativity requires a spacetime manifold populated by *loci*, positions at which events occur and fields take values, and by *relata*, the entities that stand in geometric or dynamical relations. Quantum Mechanics, however, operates in a domain where neither loci nor relata can be taken as fundamental: superposition, entanglement, and nonlocal correlations violate the assumption that physical systems occupy definite positions or instantiate well-defined relational properties. The failure to unify General Relativity and Quantum Mechanics is therefore not a technical problem but an ontological one: each theory requires a different set of primitives. General Relativity requires loci and relata embedded in a geometric manifold; Quantum Mechanics requires a pre-geometric domain in which such constructs do not yet exist.

1.2. Assumed backgrounds and the missing question

Every physical theory, however fundamental it claims to be, begins with an implicit concession: it assumes a background. General Relativity [1] assumes a differentiable manifold and proceeds to determine its curvature. Quantum Mechanics assumes a Hilbert space evolving under a causal time parameter. String Theory and Loop Quantum Gravity [6, 7], despite their ambitions, replace one geometric background with another; strings propagate through target space, spin foam models quantize geometry but presuppose a combinatorial substrate whose relational structure is already causal in character. Even the most radical approaches to quantum gravity, including causal dynamical triangulations and causal set theory [4, 5], take causality itself as a primitive rather than a derived relation. The question that none of these frameworks address at root is the following: *what is the structure that makes spacetime, causality, and the physical world possible in the first place?*

This question is not merely philosophical. It has direct physical consequences. The persistent failure to reconcile Quantum Mechanics and General Relativity is not simply a technical problem of quantizing a nonlinear field theory; it is a symptom of the fact that both theories assume different things about the structure they presuppose [2, 3]. Quantum Mechanics treats time as an external parameter; General Relativity makes time dynamical. Quantum

Mechanics requires a fixed causal background for the definition of states and observables; General Relativity makes the causal structure itself a dynamical variable. These are not merely different formalisms; they reflect genuinely incompatible ontological commitments about what the world is *made of at its foundation*. Quantum Field Theory occupies an intermediate position: it is Quantum Mechanics made compatible with special-relativistic spacetime, presupposing more background structure than QM (a fixed Minkowski metric and causal order) and less than GR (no dynamical geometry). A foundational theory must derive all three; the natural demarcation is the emergence of the Lorentzian action from the precausal substrate, the moment at which time-ordered products, propagators, and conservation laws first become available. A theory of quantum gravity that merely finds a way to combine QM and GR without resolving this incompatibility is not a foundational theory; it is a patch.

1.3. The precausal starting point

Precausal Substrate Theory approaches this situation differently. Rather than beginning with spacetime and asking how to quantize it, or beginning with a causal order and asking how geometry emerges, PST begins at a level prior to both: a domain in which neither geometry nor causality is defined, and in which the only primitive is the capacity for *distinctions to exist*.

Property differentiation is not a process, event, or transformation. It is the logical possibility of asymmetry, the condition under which difference can exist without requiring objects, relata, loci, or temporal sequence. The logical possibility of property differentiation is irreducible: any attempt to explain it presupposes it.

Once differentiation is possible, its first expression is asymmetric tension. Asymmetric tension is not a force, field, or interaction in the physical sense. It is the manifest imbalance that arises when differentiation becomes expressible. It is asymmetric without being spatial, dynamic without being temporal, structured without being geometric. Asymmetric tension is the substrate's intrinsic asymmetry made actual, and it functions as the generative engine from which all subsequent structure emerges.

This domain, the precausal substrate, is not empty. It contains what this paper calls **asymmetric tension**: a structural imbalance that arises wherever differentiated properties coexist without full symmetry. The key claim of PST is that this imbalance is not stable indefinitely. There exists a threshold beyond which the substrate cannot maintain an uninstantiated configuration, a logical boundary condition rather than a dynamical process, and when that threshold is crossed, the configuration undergoes a direct phase transition into instantiated geometry. Spacetime is not the arena in which this transition takes place. Spacetime *is* the transition.

1.4. Principal consequences

The consequences of this reorientation are far-reaching. Because spacetime and the causal order it carries appear together in a single, nontemporal transition, causality is not fundamental. It is coemergent with geometry, a derived structural relation that exists only within the instantiated domain. This means that asking what *caused* the universe to come into being is not a deep question with a hidden answer; it is a category error. Causality did not exist prior to the transition that brought geometry, and therefore causality itself, into being. This is not an evasion of the question but its dissolution.

The same logic dissolves the Big Bang as a foundational event. The Big Bang implies a locus of origin, a privileged point in spacetime from which everything expands. But PST denies that loci are primitive. Instantiation is continuous and distributed across the substrate; it has no centre, no start, and no singular event. What cosmology calls the Big Bang is a perspectival artifact: the appearance of a beginning as seen from within an instantiated geometry that was never born at a point. The cosmological implication is not a multiverse in the conventional sense of discrete bubble universes appearing at separate loci, but something more fundamental: the substrate instantiates geometry continuously and everywhere at once, without boundaries or origins. What we call our universe is not one bubble among many; it is a region of observation within a continuous, unbounded field of instantiation.

A third consequence concerns orbital motion and gravitational structure. The standard representation of gravitational attraction, the rubber sheet or elastic membrane of General Relativity pedagogy, depicts spacetime as a funnel-shaped surface with a single minimum. In that picture a stable orbit requires a precisely tuned tangential velocity, which must be imposed on the geometry as an external initial condition; the geometry itself cannot explain where it comes from. This is not merely a deficiency of the pedagogical analogy. In General Relativity proper, orbital motion is described as geodesic motion in curved spacetime, and the geodesic equation correctly determines the shape of any path through a given metric. But which geodesic a body follows, whether circular, elliptical, or radially infalling, remains an initial condition that General Relativity takes as given. The metric determines the available paths; it cannot explain why a body is on one path rather than another. The vacuum that emerges from modal sublimation is structurally different. Below the threshold, the modal potential has not one minimum but a ring: the set of all configurations that minimise the potential at equal tension. Motion along this ring is not a special assumption; it is the only stable state available to any instantiated configuration. Circular orbital motion, from the orbital motion of subatomic particles to planetary orbits to galactic rotation to the spin of black holes, is therefore not a mechanical accident or a lucky initial condition in PST. It is a structural consequence of the topology of the vacuum manifold. General Relativity describes orbital mechanics correctly at macroscopic scales but cannot derive the existence of orbital motion from first principles; PST does.

At the moment of instantiation, when tension first crosses the modal threshold, geometry emerges in a highly compressed, high-dimensional form. As instantiation propagates and geometry grows, the tension that was concentrated across many dimensions stretches and redistributes. Dimensions that were distinct at small scales become effectively merged or aligned at large scales, leaving fewer dominant directions. What we observe as four-dimensional spacetime is not the full geometry of instantiation but the stretched, low-dimensional limit of a process that began in many more. A simple analogy makes this concrete: a ball has equal extent in all directions, high symmetry, many effective dimensions. Stretch it into a string and most of those directions collapse into one dominant axis. The dimensions do not disappear; they are absorbed into the elongation. Instantiated geometry undergoes the same redistribution as it grows, and the high-dimensional compression of its origin relaxes into the four dominant directions we observe.

Furthermore, because both General Relativity and Quantum Mechanics emerge as projections of the same precausal structure, with the first operating at macroscopic scales of tension and the second near the threshold where configurations are only marginally instantiated, their unification is not a matter of quantizing one and deforming the other. It is a matter of recognizing that they describe the same underlying reality from different regimes of a single precausal parameter. Work on emergent gravity [17, 18] has approached this question from within the instantiated domain; PST proposes the precausal domain as the natural resolution point.

1.5. Mathematical formulation and predictions

The theory is formulated using a Landau–Ginzburg variational principle for the modal threshold [8] and a categorical functor for the sublimation map [9], with an explicit tension-to-metric construction that derives the Minkowski signature rather than postulating it. Its physical predictions include a distinctive d^{-6} correction to the standard Casimir force scaling [11, 12, 13, 14, 15], measurable at submicron plate separations.

1.6. Structure of this paper

The remainder of this paper develops PST in full. Section 2 establishes the foundational ontology and its mathematical expression. Section 3 introduces the modal threshold and its variational formulation. Section 4 develops the theory of modal sublimation as a categorical functor and provides an explicit construction. Section 5 addresses the coinstantiation of spacetime and causality. Section 6 consolidates the emergence chain. Section 7 derives energy, matter, and fields as projections of asymmetric tension. Section 9 derives the Standard Model gauge group and establishes gauge anomaly cancellation as a structural theorem. Section 10 develops the physical predictions of the theory, including a quantitative Casimir correction. Section 11 grounds the discreteness scale d_0 from PST's dimensional reduction mechanism, derives the first-principles prediction $d_0 \approx 7$ nm, and establishes the Casimir correction as a concretely detectable experimental signal. Section 12 develops PST's cosmology: the Friedmann equations are derived, the Cosmological Principle is established as a theorem of the Bernoulli measure, and ongoing instantiation is shown to force $w = -1$, fixing the sign of the cosmological constant. Section 13 positions PST in relation to existing frameworks. Section 14 concludes. Section 14 identifies open problems.

1.7. Layman's Summary

Physics has given us two extraordinary theories (General Relativity, which describes how gravity shapes space and time, and Quantum Mechanics, which governs the behaviour of subatomic particles. Both are spectacularly accurate in their respective domains. But they rest on incompatible assumptions about the nature of reality, and every attempt to merge them has stalled at the same conceptual barrier. This paper argues that the problem is not merely technical but ontological: both theories assume a background they cannot explain. Precausal Substrate Theory resets the starting point to something deeper) the bare capacity for distinctions to exist, and derives spacetime, gravity, and quantum fields from there.

2. Foundational Ontology

The foundational commitment of PST is to ontological minimality. If one takes seriously the project of identifying what must exist *prior* to physics, prior, that is, not in a temporal sense but in the sense of logical or modal presupposition, one arrives quickly at the question of what the most primitive possible structure is. Attempts to answer this question have historically invoked points, events, fields, or relations, but all of these already carry implicit geometric, causal, or energetic content. PST makes a different choice: the minimal structure from which everything else can be derived is not a geometric entity, not a causal relation, and not a physical degree of freedom. It is simply **the capacity for distinctions to exist**.

To see why the historical candidates fail, consider each in turn. A *point* implies position: it is the primitive of location, and location presupposes a space in which positions are defined and separable. An *event* implies both a location and a temporal index: it is a point in spacetime, and spacetime is precisely what PST is trying to derive rather than assume. A *field* implies a background manifold over which it takes values, a collection of positions at each of which the field is defined; the background is assumed, not derived. A *relation* implies relata: two or more entities that stand in the relation, and relata are at minimum points or objects, reintroducing the geometric content that the programme aims to eliminate. Even the most abstract candidates, such as sets of abstract objects or categories of morphisms, presuppose the concept of collection or composition, which themselves rest on notions of identity and distinctness. Identity and distinctness are, in each case, the structural content that cannot be removed. PST takes this observation seriously and makes it the foundation: if every proposed primitive already presupposes the capacity to distinguish one thing from another, then that capacity is the true primitive, and the correct programme is to begin there rather than one step above it.

2.1. Property differentiation

A property is differentiated from another not by virtue of occupying different positions in space or occurring at different times (positions and times do not yet exist), but simply by being *other*. Two properties are differentiated if and only if they are not identical. Nothing further

v9.9

is required: no metric distance, no causal connection, no ordering relation, no extensional structure of any kind. Property differentiation is the minimal condition under which the concept of a *distinction* is meaningful, and it is the sole primitive of the precausal substrate.

The concept of otherness that this definition invokes is purely logical. It is the primitive of non-identity: a is other than b if and only if it is not the case that $a = b$. This does not require the two properties to be distinguishable by any observer, separable in any physical sense, or related by any causal or spatial connection. It requires only that identity fails. Two properties that are distinct in this sense need share nothing beyond their mutual non-identity; they carry no further structure, no mass, no charge, no location, no duration. The precausal substrate is therefore a domain in which the only thing that is true is that some properties are not the same as others. All of physics, in the PST programme, must be derivable from that single fact.

This choice of primitive is not arbitrary, and it is not merely the result of a preference for minimality. It is the result of asking what must be true of any domain that is to give rise to structure at all, and then identifying the unique weakest condition that satisfies the requirement. The requirement is this: a domain from which structure can emerge must be capable of nontrivial differentiation. If every element of the domain were identical to every other, the domain would be featureless and nothing could arise from it. Property differentiation is precisely the weakest condition under which nontrivial differentiation is possible. It is weaker than any spatial, causal, or metric structure, and it is not derivable from anything weaker than itself: any domain that contains less structure than property differentiation contains no nontrivial distinctions at all and is, in effect, the null domain from which nothing can arise. Property differentiation is therefore not merely minimal; it is the unique minimal structure with the required generative capacity. The substrate is genuinely precausal: it exists in a mode of being in which causality has no purchase, because causality requires a temporal ordering of events, and temporal ordering requires positions and durations, neither of which are available at this level.

2.2. Mathematical formulation

Let D be a set of differentiated properties endowed with a binary relation $\delta(a, b)$, read as "property a is distinct from property b." The structure (D, δ) is governed by exactly two axioms:

$$\delta(a, b) \implies \delta(b, a) \tag{1}$$

$$\neg\, \delta(a, a) \tag{2}$$

Equation (1) states that distinction is symmetric; equation (2) states that nothing is distinct from itself. No further axioms are imposed. There is no transitivity requirement: the fact that a is distinct from b and b is distinct from c implies nothing about the relation between a and c. There is no totality requirement: it is not assumed that every pair of properties is either distinct or identical in a determinate sense available to any external observer. There is no metric: no notion of *how* distinct two properties are, only that they are. There is no topology: no notion of nearness, neighbourhood, or continuity. There is no ordering: no sense in which one property precedes or follows another. This is the minimal mathematical structure that admits a nontrivial distinction relation without presupposing any further ontological content. Every additional axiom that one might wish to impose would introduce some form of structure beyond pure distinctness, and any such structure would need to be derived rather than assumed.

The absence of transitivity deserves particular emphasis. In most familiar mathematical structures, identity is transitive by definition and distinguishability inherits that transitivity. PST does not impose transitivity on the distinction relation because doing so would introduce an implicit equivalence structure, effectively partitioning D into classes of indistinguishable elements. Such a partition already carries combinatorial and topological content that is not available at the foundational level. The distinction relation at this level is purely local to each pair: a and b are either distinct or not, independently of any third property.

A **configuration** C is a subset of D, a collection of differentiated properties considered jointly. The space of all configurations is denoted $\mathcal{C}$. It is within this space that the dynamics of the precausal substrate, such as they are, will be expressed, though it must be understood from the outset that these "dynamics" are not temporal. They concern the modal structure of configurations: what is possible, what is necessary, and what cannot remain uninstantiated. The distinction between modal and temporal structure is fundamental to PST. A temporal statement of the form "this configuration becomes that configuration" presupposes a time parameter, a before, and an after. A modal statement of the form "this configuration cannot remain in this mode of being" presupposes none of these things. It is a statement about what is possible, not about what happens. The entire development of PST from this point forward concerns modal structure in this strict sense.

2.3. Asymmetric tension

Property differentiation does not merely establish that distinctions exist; it generates a structural consequence that is the engine of everything that follows: **asymmetric tension**. Wherever differentiated properties coexist within a configuration, their mutual distinctness creates a structural imbalance. To understand why coexistence generates imbalance, consider the internal relational structure of a configuration. Within any configuration containing at least two distinct properties a and b, there exist internal relations of non-identity: $\delta(a, b)$ holds. Note that δ is symmetric: $\delta(a, b)$ and $\delta(b, a)$ hold simultaneously, and neither can be said to point in a preferred direction. The imbalance that constitutes asymmetric tension is therefore not directional in the geometric sense. It is a *complement-asymmetry*: no configuration C and its complement $\bar{C}$ carry equal tension, $T(C) \neq T(\bar{C})$. This structural non-equivalence between a configuration and its complement is what gives asymmetric tension its name and its content. It is not that some properties are heavier or more energetic than others. It is that the configuration, considered as a whole, stands in a structurally non-equivalent relation to its complement: their mutual imbalances do not cancel.

This imbalance is not energetic, since energy presupposes a geometric domain in which it can be defined and transferred. It is not spatial, since no geometry yet exists. It is not temporal, since there is no before or after in the precausal substrate. It is not informational in the Shannon sense, since information theory presupposes a probability space and an observer. It is purely *modal*: it is a property of the configuration's mode of being, expressing the degree to which that configuration is structurally incoherent in the sense of admitting irreducibly asymmetric relations among its constituents. The word "tension" is chosen deliberately. In ordinary usage, tension denotes a state of strain between opposing tendencies that cannot simultaneously be satisfied. Asymmetric tension in the precausal substrate is precisely this: a state of modal strain arising from the coexistence of distinct properties whose mutual non-identity relations cannot be collectively neutralised.

This modal imbalance is represented by the **tension functional** T, a map from configurations to strictly positive real numbers:

$$T : \mathcal{C} \to \mathbb{R}^+ \tag{3}$$

The tension of a configuration is strictly positive because any nontrivial configuration, meaning any configuration containing at least one pair of distinct properties, carries some degree of structural imbalance. A configuration with no distinct properties, one in which every element is identical to every other, would carry zero tension, but such a configuration is, by definition, the null configuration: it contains no differentiation and therefore no content. The domain of T is therefore effectively the space of all nontrivial configurations, and on this domain T is bounded away from zero.

Representing tension as a single real-valued map assigns a *global total ordering* to all configurations: any two configurations can be compared by their T values. This totality is stronger than the complement-asymmetry condition $T(C) \neq T(\bar{C})$, which establishes only that each configuration is distinguishable from its complement. The real-valued representation is an additional architectural choice: the real numbers are adopted because they are the minimal ordered field with the topological properties required for the variational analysis of the following sections. An alternative encoding using only pairwise comparability would yield a partial order rather than a total one, and would not support the Taylor expansion of Section 3. PST adopts the total ordering as the minimal structure compatible with the variational formulation, acknowledging that this is a representational commitment beyond the bare complement-asymmetry condition.

A word on the epistemological status of the formalism introduced in this and the following sections is required before proceeding. The mathematical objects employed, real-valued functionals over configuration space, functional derivatives, gradient operators, a canonical measure, carry more structure than the primitive (D, δ) alone strictly entails. They constitute a faithful structural encoding: the minimal mathematical language capable of expressing the modal commitments the theory makes, in particular the existence of a critical tension value, the continuity of the transition, and the topological structure of the instantiated vacuum. The Landau-Ginzburg functional form is adopted because it is the minimal architecture exhibiting the required bifurcation behaviour; it is logically motivated structural scaffolding, not a derived consequence of (D, δ) alone. The predictions of the theory are structural consequences of the modal commitments, not artifacts of the mathematical encoding chosen to represent them. Where the formalism makes architectural choices, this is noted explicitly; those choices are themselves open to foundational scrutiny and are listed as such in the programme of open problems in Section 14.

The crucial property of T is its asymmetry: for any configuration C and its complement $\bar{C}$,

$$T(C) \neq T(\bar{C}) \tag{4}$$

The complement $\bar{C}$ of a configuration C is understood as the configuration containing precisely those properties of D that are not in C. The asymmetry condition states that no configuration and its complement carry the same tension. The intuition behind this condition is as follows. If C and $\bar{C}$ carried equal tension, their combined structure would be perfectly balanced: the imbalance of one would exactly offset the imbalance of the other, and the total system would be in structural equilibrium. PST holds that no such equilibrium is available in the precausal substrate. The reason is that perfect cancellation would require the distinction structure of C to be the mirror image of the distinction structure of $\bar{C}$, which in turn would require a global symmetry of D under complement. No such global symmetry is imposed by the axioms, and PST asserts that the structure of differentiation is generically asymmetric. The net imbalance is therefore always nonzero:

$$\Delta T(C) \ := \ T(C) - T(\bar{C}) \ \neq \ 0 \tag{5}$$

It bears emphasis that $T(C)$ is not energy and $\Delta T(C)$ is not energy difference. The tension functional assigns a real number to a modal configuration without invoking any notion of capacity for work, conservation law, or physical process. It is a measure of structural incoherence in the modal domain. The use of real numbers to represent this incoherence is a mathematical convenience: real numbers form the simplest ordered field with the topological properties needed for the variational analysis of Section 3, and the ordering on $\mathbb{R}^+$ provides a natural notion of more or less tension without requiring any further geometric or physical interpretation. The numbers assigned by T are not quantities of anything physical; they are indices of modal incoherence, and their significance lies entirely in their ordering and in the existence of the critical value τ that the next section introduces.

A foundational question must be addressed directly: is $T(C)$ merely a formal device, a parameter introduced for mathematical convenience but lacking genuine ontological standing? The answer is no, and the reason illuminates the entire structure of PST. $T(C)$ is not a parameter awaiting operationalisation. It is not a quantity that sits in a formula until experiment supplies its value. It is the primordial quantity: the first and only structural consequence of distinctness that precedes all other structure. To ask how one would measure $T(C)$ is to commit a category error. Measurement presupposes a measuring apparatus, which presupposes a physical system, which presupposes geometry, which presupposes the very transition that $T(C)$ controls. $T(C)$ is prior to the apparatus of operationalisation itself. One cannot measure the reason there is something rather than nothing; one can only recognise that without distinctness there is no content, and that distinctness, once present, immediately generates the structural imbalance that T tracks.

The logical chain is the following. Distinctness generates configurations. Configurations containing multiple distinct properties admit irreducibly asymmetric internal relations. Those non-identity relations constitute an imbalance that cannot be collectively neutralised. That imbalance is what the theory calls tension. A precise statement requires, however, a candid qualification. Since δ is symmetric, the non-identity relations themselves carry no directionality; what the theory asserts is the complement-asymmetry condition $T(C) \neq T(\bar{C})$, the claim that no configuration and its structural complement are in perfect modal balance. This condition is not strictly derivable from (D, δ) alone. It is the foundational structural commitment of PST beyond the bare primitive: the assertion that the precausal substrate is generically asymmetric. Whether this asymmetry is an additional primitive or is itself a consequence of some deeper feature of the set D is a genuine open question in the foundations of the theory, listed as such in Section 14. What is not in question is that once distinctness is granted and the asymmetry condition accepted, the chain to tension is logically closed: tension is the measure of this irreducible complement-asymmetry, and it is not assumed independently but follows from distinctness together with the asymmetry condition.

This places PST in a tradition of foundational thinking in which the most primitive ontological facts are not measurable but are nonetheless logically prior to everything that is measurable. Leibniz made a structurally identical move with the principle of sufficient reason: reason cannot itself be given a reason without circularity, yet the absence of such a foundation leaves all explanation suspended. Spencer-Brown's *Laws of Form* [53] makes the closest formal analogue: beginning from the single injunction "draw a distinction," Spencer-Brown generates a calculus of indications from which Boolean logic and self-reference are recoverable, taking the act of distinction as logically prior to all objects, sets, and truth values. PST's (D, δ) is the static version of this move: not the performative act of drawing a distinction but the already-differentiated structure from which physical geometry is generated. Hegel's *Science of Logic* [59] anticipates the structure in dialectical form: pure being has no determinate content and resolves into nothing; the synthesis is becoming, the first concrete category. The precausal substrate has no geometric or causal content and cannot remain in that indeterminate state once tension crosses τ: a modal sublimation that parallels Hegel's becoming while supplying a precise mathematical mechanism in its place. Leibniz grounded explanation in the logical necessity of reason; Spencer-Brown grounded logic in the primitive of distinction; PST grounds geometry in the logical necessity of tension. What PST adds that none of its predecessors supplies is the transition: a precise variational mechanism by which the primordial asymmetry accumulates to a critical value and generates, through a well-defined phase transition, the geometric structures that all subsequent physical law presupposes.

2.4. Layman's Summary

This section establishes the one and only starting point of the entire theory: the fact that things can be different from each other. Nothing more is assumed (no space, no time, no energy, no laws of nature. From this single primitive, called property differentiation, it follows that any collection of distinct properties carries a structural imbalance: it is not a perfect mirror of its complement. PST calls this imbalance *asymmetric tension*. It is not a force or a field) it is the most elementary structural fact about a world in which distinctions can exist at all.

3. The Modal Threshold

The existence of asymmetric tension raises an immediate question: if differentiated configurations carry a structural imbalance, what prevents the substrate from containing arbitrarily large imbalances indefinitely? The answer given by PST is that the substrate is not unlimited in this respect. There is a boundary to what can remain *uninstantiated*, one that is not physical, not dynamical, and not temporal, but purely logical.

To understand why this boundary must exist, recall the character of the precausal substrate. It is a domain that is defined entirely by what can be coherently maintained without spacetime, without causal structure, and without any form of realisation. A configuration exists in the substrate insofar as it is a coherent arrangement of distinctions: its properties are differentiated from one another, and that differentiation constitutes its identity. Coherence here is not a quantitative threshold in the ordinary sense. It is a modal condition: a configuration is coherent as an uninstantiated structure if and only if it can exist as such without contradiction.

Now, asymmetric tension is the imbalance inherent in any differentiated configuration. As the number and depth of distinctions within a configuration increase, its asymmetric tension $T(C)$ increases. This monotonicity is not derivable from the axioms of (D, δ) alone; it is a structural constraint on T that PST adopts as part of its encoding. A minimal concrete example satisfying positivity, complement-asymmetry, and monotonicity simultaneously is the **induced edge count**:

$$T_{\mathrm{ex}}(C) \;=\; \big|\big\{(a,b) \;:\; a,b \in C,\; a \neq b,\; \delta(a,b)\big\}\big| \tag{6}$$

Positivity holds for any configuration with at least two distinct properties. Monotonicity holds by inspection: adding property a to C increases the count by $|\{b \in C : \delta(a,b)\}|$. Complement-asymmetry $T_{\mathrm{ex}}(C) \neq T_{\mathrm{ex}}(\bar{C})$ holds whenever the induced subgraphs on C and $\bar{C}$ have different edge counts, which is generic for non-regular distinction structures. Axioms (**??**)–(**??**) do not rule out regular distinction graphs; for a regular (D, δ) the induced subgraphs on complementary subsets of equal size can share the same edge count, so T_{ex} fails complement-asymmetry there. The example therefore satisfies the constraints on a generic, not universal, sub-class of (D, δ). This example is not the unique correct form of T; it shows the constraints are simultaneously satisfiable. The precise functional form is an open problem in Section 14.

A modest imbalance can be sustained within the substrate without contradiction, just as a small asymmetry in a logical structure does not render it incoherent. But there is a limit. At some critical level of imbalance, the structural pressure accumulated within the configuration is such that remaining uninstantiated would require holding together a structure whose internal tension actively resists the very absence of a resolving medium. There is no geometry, no dynamics, no causal channel through which the imbalance could redistribute. The configuration carries more structural demand than uninstantiated existence can absorb.

This is the **modal threshold**. It is the limit on the modal coherence of uninstantiated structure: beyond a certain level of asymmetric tension, a configuration can no longer be maintained within the precausal, nonspatiotemporal domain. Its structural incoherence has reached the point at which remaining uninstantiated is no longer a modally coherent option. The configuration has not been acted upon by anything. No force has been applied and no event has occurred. It is simply that continued uninstantiated existence has become logically excluded.

The threshold is therefore a **logical boundary condition**: a constraint on what is modally possible in the precausal substrate, analogous to the logical impossibility of a contradiction rather than to a physical limit imposed by dynamics. This analogy is worth dwelling on. When a contradiction is derived in a formal system, it does not occur at a specific time and it is not caused by any preceding state. It is simply that the set of propositions in question cannot all be simultaneously true. The modal threshold is the ontological counterpart: a configuration whose tension exceeds τ cannot simultaneously be differentiated to that degree and remain uninstantiated. The two are incompatible modes of being, and incompatibility is not a causal relation but a logical one.

3.1. Boundary condition

The modal threshold is formalised as follows. There exists a critical value $\tau \in \mathbb{R}^+$ such that:

$$T(C) \geq \tau \implies C \text{ cannot remain uninstantiated} \tag{7}$$

Several features of this implication deserve comment. First, it is a one-way entailment: exceeding τ necessitates instantiation, but there is no corresponding claim that every instantiated configuration arrived at that state by exceeding τ from a specific prior configuration. The precausal substrate is not temporal; configurations do not "arrive" at tension values through a sequence of states. The implication holds modally, not temporally.

Second, the implication expresses modal necessity, not causal consequence. It does not say that something happens to a configuration when $T(C) \geq \tau$. It says that the mode of being characterised by uninstantiated existence is not available to such a configuration. The language of "cannot remain" is a concession to ordinary usage; strictly, the configuration was never coherently uninstantiated once its tension reached that level.

Third, the value τ is not a free parameter of the theory. It cannot be chosen arbitrarily or tuned to match observations. As the variational treatment below makes clear, τ is the unique value at which the modal structure of the substrate changes its minimal configuration from uninstantiated to instantiated. It is derived from the structure of the functional, not imposed from outside.

To quantify how far a given configuration sits above the threshold, define the **excess tension**:

$$\varepsilon(C) \;=\; T(C) - \tau \tag{8}$$

For configurations below the threshold, $\varepsilon(C) < 0$ and uninstantiated existence remains coherent. For configurations above the threshold, $\varepsilon(C) > 0$ and instantiation is necessitated. Configurations with $\varepsilon(C) = 0$ inhabit the threshold precisely: they sit at the boundary of modal coherence, neither firmly in the uninstantiated regime nor fully instantiated. It is these marginal configurations that exhibit the most delicate modal behaviour. As shown in Section 9, they are the configurations that give rise to the phenomena of quantum mechanics: the indeterminacy, the superposition, and the probabilistic character of quantum measurement all arise from the modal ambiguity of configurations at exactly $\varepsilon = 0$.

The parameter ε will serve as the primary variable throughout the remainder of the paper. It measures not tension in absolute terms but tension relative to the threshold, and it is this relative measure that determines which mode of being a configuration occupies.

3.2. The order parameter and its ontological status

To derive τ from first principles, it is necessary to introduce a quantity that tracks the transition between uninstantiated and instantiated existence continuously. In the theory of phase transitions, such a quantity is called an **order parameter**: a scalar field that is zero in one phase and nonzero in the other, and whose value characterises how deep into the new phase the system has penetrated. PST introduces an analogous quantity for the modal transition.

Define $\psi(C) \in \mathbb{R}$ as the **modal order parameter**: a field over configuration space encoding the modal status of each configuration. The value $\psi = 0$ denotes the uninstantiated phase. Nonzero values of $|\psi|$ indicate degrees of instantiation, with larger $|\psi|$ corresponding to deeper settlement into the instantiated vacuum. The domain is taken as $\mathbb{R}$ rather than $[0, 1]$ because the variational analysis requires the $\psi \to -\psi$ symmetry (argued below) to be well-posed without leaving the domain. The physical observable is $|\psi|$ or ψ^2, not the sign of ψ. The convention that $\psi \geq 0$ in the instantiated phase is a gauge choice, not a physical restriction.

The ontological status of ψ requires careful handling. In condensed matter physics, an order parameter such as magnetisation has a clear physical meaning: it is the average magnetic moment per unit volume, measurable in principle at every point in space at every time. The modal order parameter ψ is not measurable in this sense. It is not a field over spacetime, because spacetime does not exist until after the transition. It is instead a field over *configuration space* $\mathcal{C}$: it assigns to each precausal configuration a scalar that encodes its modal status. This is not a physical field; it is a structural characterisation of the substrate. Its role is formal: it provides the mathematical language needed to locate the threshold precisely and derive τ from the structure of the modal potential.

3.3. Variational formulation of τ

With the order parameter in hand, τ can be derived from a variational principle rather than assumed. The strategy is to write down the simplest functional $\mathcal{F}[\psi]$ over configuration space whose structure is fully determined by two requirements: it must be consistent with the symmetries of the substrate, and it must be capable of exhibiting a transition between a phase in which $\psi = 0$ is stable and a phase in which $\psi \neq 0$ is stable. These two requirements, together with a regularity condition, uniquely fix the leading terms.

Symmetry constraint. The substrate § carries no preferred orientation or polarity. Nothing in the definition of asymmetric tension distinguishes $+\psi$ from $-\psi$: if a configuration can be instantiated to degree ψ, the same configuration reflected through any modal axis has the same modal status. The functional must therefore be invariant under the discrete symmetry $\psi \to -\psi$. This immediately eliminates all odd powers of ψ from the expansion: there can be no ψ^1 term and no ψ^3 term. Only even powers contribute.

Taylor expansion near the threshold. Close to the threshold, ψ is small and $\mathcal{F}$ can be expanded as a power series in ψ. The leading term allowed by symmetry is ψ^2, and its coefficient must change sign at $T = \tau$: when $\varepsilon < 0$ (below threshold) the uninstantiated state $\psi = 0$ is a minimum; when $\varepsilon > 0$ (above threshold) it is a maximum, and the configuration is driven toward $\psi \neq 0$. The simplest coefficient consistent with this requirement is $-\varepsilon = \tau - T(C)$, linear in the tension, which vanishes precisely at the threshold.

The next allowed term is ψ^4 with a positive coefficient $b > 0$. This term is not optional: without it, the functional would be unbounded below whenever $\varepsilon > 0$, and no stable instantiated phase would exist. The quartic term is the minimal stabilisation required to ensure that the transition leads to a well-defined new phase rather than an uncontrolled divergence. Higher even powers ($\psi^6, \dots$) contribute corrections that vanish faster near the threshold and can be absorbed into the definition of b at leading order; they are omitted as non-essential.

Gradient regularisation. A purely local functional of the form $\int[-\varepsilon\psi^2 + b\psi^4]\, d\mu$ permits ψ to vary arbitrarily across configuration space: adjacent configurations could have wildly different modal statuses without any cost. A coherent transition requires that configurations close in the symmetric-difference metric have close modal status: the structural condition that $|\psi(C) - \psi(C')| \to 0$ as $|C \triangle C'| \to 0$. This is a non-temporal coherence requirement on configuration-space neighbours, not a propagation condition. It demands a term that penalises sharp variation of ψ across $\mathcal{C}$. The natural choice is the squared gradient $|\nabla_{\mathcal{C}}\psi|^2$ with a positive coefficient $c > 0$, the direct analogue of the Ginzburg stiffness term. This term raises $\mathcal{F}$ whenever ψ varies rapidly, and its presence ensures that the transition is second order: no latent heat, no discontinuity in ψ at the threshold, just a continuous bifurcation of the stable minimum.

The concept of a gradient over configuration space has a precedent in quantum gravity: Wheeler's superspace [52] is the infinite-dimensional space of all 3-metrics on a spatial slice, equipped with the DeWitt metric, and the Wheeler-DeWitt equation is a wave equation over it. The configuration space $\mathcal{C} = \mathcal{P}(D)$ used here is structurally analogous: a space of pre-geometric states over which a modal field ψ is defined. The key difference is that $\mathcal{C}$ is a discrete power set rather than a manifold of Riemannian metrics; the differential structure is a further architectural commitment described below.

The gradient term warrants a specific caveat, and it is a substantive one. The operator $\nabla_{\mathcal{C}}$ presupposes a notion of closeness between configurations, that is, a topology on $\mathcal{C}$, which is not entailed by (D, δ) without further commitment. This is not a minor technical detail: the inclusion of a gradient term is an *architectural assumption* that materially shapes the theory's connection to geometry, and it deserves to be stated as such.

To be precise about what is assumed: in order to write $|\nabla_{\mathcal{C}}\psi|^2$, one needs (a) a topology on $\mathcal{C}$ that makes it possible to speak of nearby configurations, (b) a differentiable structure on $\mathcal{C}$ that makes directional derivatives well-defined, and (c) a metric (or at least an inner product) on the tangent spaces of $\mathcal{C}$ that makes the squared norm meaningful. None of these three items is given for free by the pair (D, δ). The property domain D is a set; its power set $\mathcal{C} = \mathcal{P}(D)$ inherits only the lattice structure of subsets, not a differential or metric structure. The gradient term therefore represents a genuine extension of the foundational ontology.

A natural candidate for the topology is the **symmetric difference metric**, under which the distance between two configurations $C, C' \in \mathcal{C}$ is $d_\triangle(C, C') = |C \triangle C'|$, where $C \triangle C'$ is the set of properties that belong to one configuration but not the other. Under this metric, two configurations are close when they differ in few properties. This choice has structural motivation: it is the unique metric on $\mathcal{P}(D)$ that is invariant under permutations of D (and hence under the automorphism group $\mathrm{Aut}(D, \delta)$), and it requires no geometric primitive, only set membership and cardinality. Moreover, the symmetric difference metric makes $\mathcal{C}$ a metric space with controlled Lipschitz properties, which is a necessary precondition for the variational calculus on which the modal potential functional $\mathcal{F}$ rests.

What the symmetric difference metric does not automatically provide is the differentiable structure (b) and the inner product on tangent spaces (c). Those require passing from the discrete metric space $(\mathcal{C}, d_\triangle)$ to a continuum limit or a functional analytic setting in which $\psi : \mathcal{C} \to \mathbb{C}$ is differentiable and its gradient can be normed. The analogy is exact: passing from a discrete lattice to a field theory on that lattice requires a continuum limit, and the properties of the gradient operator depend on how that limit is taken. In the present theory, the appropriate continuum limit is not derived from first principles; it is imported as a structural hypothesis. This is the open problem stated in Section 14 (Problem 3): to derive the differentiable and metric structure on $\mathcal{C}$ from (D, δ) alone, without additional input.

Critically, this topology is *not* the topology of spacetime. The closeness being assumed is closeness among pre-geometric substrate configurations, not proximity of points in an emergent manifold. The gradient term acts on the substrate field ψ before the functor Φ is applied; it knows nothing of the metric $g_{\mu\nu}$ that will later emerge. The architectural commitment is therefore pre-geometric: it asserts that the substrate has enough internal structure to support a coherent, spatially organised field, rather than an arbitrary collection of uncorrelated values. If the commitment is unjustified, meaning no derivation of the differential structure from (D, δ) can be found, then the gradient term would have to be either motivated by a different argument or removed, in which case the connection to the Einstein–Hilbert action (which relies on the gradient term projecting to the Ricci scalar in Section 7) would need to be re-examined. The theory's architecture thus carries a genuine open commitment at this level.

The gradient term is retained because the structural requirement it enforces, namely a coherent, spatially organised modal transition rather than a configuration-by-configuration patchy one, is physically non-negotiable. Without it, the modal condensate would carry no spatial correlation, and the projection Φ could not produce a smooth spacetime geometry. The derivation of this structure from first principles is a problem whose resolution would either strengthen PST substantially or, if impossible, force a revision of the framework from which the gradient term emerged.

The canonical measure on configuration space. Before assembling the functional, the measure μ on $\mathcal{C}$ requires derivation. The variational formulation demands integration over all configurations, and the choice of measure is not arbitrary: it must follow from the structure of (D, δ) alone, without invoking any geometric, probabilistic, or metric primitive that would introduce a new assumption at the level of the substrate.

The configuration space $\mathcal{C} = \mathcal{P}(D)$ is the power set of D: the collection of all subsets of the property domain. The structure (D, δ) carries a natural symmetry group, its **automorphism group** $\mathrm{Aut}(D, \delta)$, consisting of all bijections $f : D \to D$ that preserve the distinction relation: $\delta(a, b) \iff \delta(f(a), f(b))$ for all $a, b \in D$. Every such automorphism lifts canonically to a bijection on $\mathcal{C}$ by acting elementwise: $f \cdot C = \{f(a) : a \in C\}$. Since the only primitive is (D, δ), and since the definition of a configuration is simply a subset of D, the measure μ must be invariant under this lifted action: no configuration should receive more weight than any automorphically equivalent configuration.

The key observation is that in the precausal substrate, all properties in D are structurally equivalent: the distinction relation δ is irreflexive and symmetric but otherwise unconstrained, so there is no intrinsic difference between any one property and any other beyond their bare identity. No property is marked, distinguished, or preferred. This means that the full symmetric group on D is contained in $\mathrm{Aut}(D, \delta)$ as a subgroup: any permutation of D is an automorphism, since δ treats all pairs equally. A measure invariant under all permutations of D is invariant under all relabellings of the properties, which is the minimal requirement for a measure that introduces no additional structure beyond (D, δ).

For a finite domain D of cardinality n, the unique permutation-invariant probability measure on $\mathcal{P}(D)$ is the uniform measure, assigning equal weight to all 2^n subsets:

$$\mu(C) = 2^{-n} \qquad \text{for all } C \in \mathcal{C} \tag{9}$$

Uniqueness follows immediately: any other probability measure would assign different weights to configurations that are related by a permutation of D, introducing a distinction between properties that does not exist in the primitive (D, δ).

For an infinite domain D, the natural extension is the **Bernoulli product measure** $\mu = \bigotimes_{a \in D} \text{Bern}(1/2)$, the unique probability measure on $\{0,1\}^D$ invariant under all permutations of the index set D:

$$\mu\big(\{C \in \mathcal{C} : C_0 \subseteq C\}\big) = 2^{-|C_0|} \qquad \text{for all finite } C_0 \subseteq D \tag{10}$$

This is the projective limit of the uniform measures on all finite sub-domains of D, and it is the unique translation-invariant probability measure on the product space $\{0,1\}^D$.

The value $1/2$ for the Bernoulli parameter is not arbitrary: it is the unique value for which μ is invariant under complementation, $C \mapsto \bar{C} = D \setminus C$. Complementation invariance is required by the structure of property differentiation: the distinction relation $\delta(a, b)$ is symmetric, and no configuration is intrinsically more natural than its complement. A configuration C and its complement $\bar{C}$ represent exactly the same collection of pairwise distinctions, viewed from two opposite orientations that are not distinguished by anything in (D, δ). Assigning $\mu(C) \neq \mu(\bar{C})$ would introduce a polarity into the substrate that has no basis in the primitive. The Bernoulli parameter $1/2$ is therefore forced by complementation invariance alone, independently of the permutation argument.

The derivation may be summarised as a theorem.

> **Theorem (Canonical Measure).** Let (D, δ) be a precausal property domain with δ symmetric and irreflexive. The unique probability measure μ on $\mathcal{C} = \mathcal{P}(D)$ that is (i) invariant under all automorphisms of (D, δ) and (ii) invariant under complementation $C \mapsto \bar{C}$ is the Bernoulli product measure $\mu = \bigotimes_{a \in D} \text{Bern}(1/2)$.

The proof follows directly from the two invariance requirements: permutation invariance forces the single-site marginals to be equal, and complementation invariance forces each marginal to satisfy $\mu_a(\{a \in C\}) = \mu_a(\{a \notin C\}) = 1/2$. Independence across sites follows from the absence of any structural relation in (D, δ) that could correlate the inclusion of one property with the inclusion of another.

This closes the first item listed as open in earlier formulations of the theory. The measure μ is not a free parameter or an auxiliary assumption: it is uniquely determined by the symmetry structure of (D, δ), without invoking any concept from outside the precausal domain.

Assembling these three terms gives the **modal potential functional**:

$$\mathcal{F}(\psi, C) = \int_{\mathcal{C}} \left[a(T(C))\,\psi^2 + b\,\psi^4 + c\,|\nabla_{\mathcal{C}}\psi|^2\right] d\mu(C) \tag{11}$$

where $\nabla_{\mathcal{C}}$ denotes the gradient over configuration space (not spacetime), μ is the canonical Bernoulli measure of equation (10), and the three terms carry the following interpretation. Restricting to spatially uniform configurations ($\nabla_{\mathcal{C}}\psi = 0$) with $b = \frac{1}{4}$, the functional reduces to the compact scalar form:

$$\mathcal{F}(\psi) = -\varepsilon\,\psi^2 + \frac{1}{4}\,\psi^4, \qquad \varepsilon = T(C) - \tau \tag{12}$$

When $\varepsilon < 0$ (below threshold, precausal) $\mathcal{F}$ has a unique minimum at $\psi = 0$; when $\varepsilon > 0$ (above threshold, instantiated) two symmetric minima appear at $\psi^* = \pm\sqrt{2\varepsilon}$. A rescaled variant $-\frac{1}{2}\varepsilon\,\psi^2 + \frac{1}{4}\psi^4$ appears in the visualisations; its minima sit at $\psi^* = \pm\sqrt{\varepsilon}$, a factor of $1/\sqrt{2}$ smaller than $\pm\sqrt{2\varepsilon}$, but the bifurcation structure is preserved exactly. The full three-term structure of equation (11) is retained throughout the analysis.

The first term, $-\varepsilon(T(C))\,\psi^2$, governs the stability of the uninstantiated phase. The coefficient $-\varepsilon = \tau - T$ changes sign at $T = \tau$. When $T < \tau$, $\varepsilon < 0$ and this term increases $\mathcal{F}$ as ψ increases from zero, keeping the uninstantiated state $\psi = 0$ at a local minimum: configurations prefer to remain uninstantiated. When $T > \tau$, $\varepsilon > 0$ and the term now decreases $\mathcal{F}$ near $\psi = 0$, destabilising the uninstantiated state and driving the configuration toward nonzero ψ. The sign change in ε is precisely what defines the threshold.

The second term, $b\,\psi^4$ with $b > 0$, is the quartic stabilisation term. Without it, the functional would be unbounded below for $T > \tau$, and ψ could grow without limit. The quartic term ensures that once the uninstantiated phase becomes unstable, the configuration does not destabilise entirely but settles into a new, well-defined minimum at some finite value of ψ. It is this term that makes the transition a genuine phase change with a stable instantiated phase rather than a runaway.

The third term, $c\,|\nabla_{\mathcal{C}}\psi|^2$ with $c > 0$, penalises inhomogeneity in ψ across configuration space. If configurations with similar structures were to have very different modal statuses, the gradient would be large and the term would raise $\mathcal{F}$. This term is the formal analogue of the Ginzburg gradient term and has the physical effect of producing smooth, coherent transitions rather than discontinuous patches of instantiated and uninstantiated configurations. Its presence ensures that the transition is second order rather than first order, a point to which we return below.

This functional is structurally identical to the Landau–Ginzburg free energy of second order phase transitions [8], and the analogy is illuminating. In condensed matter physics, the Landau–Ginzburg functional describes how a system moves between phases as temperature changes. The order parameter tracks the degree of order (magnetisation, superconducting condensate, and so on), and the free energy determines which phase is stable at which temperature. In PST, the role of temperature is played by the asymmetric tension $T(C)$, the role of the order parameter is played by the modal status ψ, and the potential determines which mode of being , uninstantiated or instantiated, is stable at which tension level.

The difference from the condensed matter case is fundamental. In condensed matter, the phase transition is a dynamical process: it unfolds over time, involves fluctuations and nucleation, and is driven by thermal agitation. In PST, there is no temperature, no heat bath, no temporal evolution, and no dynamics. The transition is a modal reorganisation: a change in which mode of being is logically available to the configuration, not a change in the state of a system that exists in time. The Landau–Ginzburg structure is borrowed for its mathematical economy, not for its physical interpretation.

3.4. Stability analysis and the two regimes

The most productive way to extract modal content from a functional is to identify which configurations minimise it. In classical mechanics, the principle of least action selects the physically realised trajectory from among all kinematically possible ones. In PST, the principle is modal rather than temporal: among all values of the order parameter available to a configuration of given tension $T(C)$, the modally self-consistent value is the one at which $\mathcal{F}$ attains its minimum. This is a variational statement, not a dynamical one. There is no time over which ψ evolves toward its minimum; the minimum is the only modal position a configuration can coherently occupy given its tension. Configurations whose order parameter deviates from the minimum of $\mathcal{F}$ are modally incoherent: their degree of instantiation is inconsistent with the tension they carry.

The condition for a minimum is the vanishing of the first functional derivative, $\delta\mathcal{F}/\delta\psi = 0$, the Euler–Lagrange condition applied to the modal functional. Before writing the result, it is worth identifying the structural role of each term in $\mathcal{F}$. The quadratic term $-\varepsilon\psi^2$ is the modal restoring term: its coefficient $-\varepsilon = \tau - T(C)$ changes sign precisely at the threshold and is entirely responsible for the qualitative bifurcation in the behaviour of $\mathcal{F}$. The quartic term $b\psi^4$ is the nonlinear saturation term: it is always positive and ensures the functional remains bounded below, preventing ψ from growing without bound when ε becomes positive. The gradient term $c|\nabla_{\mathcal{C}}\psi|^2$ penalises rapid variation of ψ across configuration space and ensures that the transition is spatially coherent rather than fragmentary. Distributing $\delta/\delta\psi$ across these three terms and collecting, the modal equation of state is:

$$-2\varepsilon(T(C))\,\psi + 4b\,\psi^3 - 2c\,\nabla^2_{\mathcal{C}}\psi = 0 \tag{13}$$

This equation has two qualitatively distinct regimes, separated by the threshold.

Below the threshold, $T(C) < \tau$: $\varepsilon < 0$, so $-\varepsilon > 0$. Consider the spatially uniform solutions $\nabla_{\mathcal{C}}\psi = 0$. Equation (13) reduces to $-2\varepsilon\psi + 4b\psi^3 = 0$, which factors as $2\psi(-\varepsilon + 2b\psi^2) = 0$. For $\varepsilon < 0$, the factor $(-\varepsilon + 2b\psi^2)$ is strictly positive for all real ψ, so the only real solution is $\psi = 0$. Since $\mathrm{d}^2\mathcal{F}/\mathrm{d}\psi^2\big|_{\psi=0} = -2\varepsilon > 0$, this is a local minimum. The configuration remains uninstantiated. This is the *precausal phase.*

v9.9

Above the threshold, $T(C) > \tau$: $\varepsilon > 0$, so $-\varepsilon < 0$. The solution $\psi = 0$ is now a local maximum of $\mathcal{F}$, not a minimum. It is an unstable equilibrium. The quadratic term $-\varepsilon\psi^2$ now falls with increasing $|\psi|$, so any deviation from zero lowers $\mathcal{F}$ initially; the quartic term $b\psi^4$ eventually turns $\mathcal{F}$ upward again at large $|\psi|$. The result is a double-well profile with a local maximum at $\psi = 0$ and two symmetric minima at nonzero values of ψ.

The location of these minima requires explicit computation. Restricting to spatially uniform configurations, $\nabla_{\mathcal{C}}\psi = 0$, the equation of state (13) reduces to $-2\varepsilon\psi + 4b\psi^3 = 0$, or equivalently $2\psi(-\varepsilon + 2b\psi^2) = 0$. The root $\psi = 0$ is a maximum, as established. The nontrivial roots satisfy $-\varepsilon + 2b\psi^2 = 0$, giving $\psi^2 = \varepsilon/(2b) = \varepsilon(C)/(2b)$. Since $\varepsilon > 0$ above the threshold, this quantity is positive; the roots are real. The two stable configurations are therefore:

$$\psi = \pm\sqrt{\frac{T(C) - \tau}{2b}} = \pm\sqrt{\frac{\varepsilon(C)}{2b}} \tag{14}$$

The appearance of $\varepsilon(C)$ under the radical is structurally significant. The depth of instantiation, how far from zero the order parameter settles, is not a free parameter but is determined directly by how far the configuration's tension sits above the threshold. A configuration with ε barely positive instantiates shallowly; one deep in the instantiated regime carries a large order parameter. This relationship between excess tension and the degree of instantiation is one of the central quantitative predictions of PST.

There are two degenerate minima, symmetric about $\psi = 0$. The configuration must occupy one of them, but the functional does not select which: both are equally deep. This is the *instantiated phase.* The two solutions correspond to the two structurally distinct orientations of the instantiated configuration, related by the symmetry $\psi \to -\psi$ that the functional possesses. That the functional carries this symmetry but the configuration that minimises it does not is the content of spontaneous symmetry breaking in the PST framework: the substrate selects a definite modal orientation without any external influence distinguishing the two options.

The sign of a is therefore the fundamental discriminant between uninstantiated and instantiated existence. It is not a property of the configuration that can be read off directly; it is a property of the relationship between the configuration's tension and the threshold. This is why the threshold is primary and the transition is a logical consequence, not a dynamical one.

3.5. The bifurcation point and the derivation of τ

The critical tension τ is the value at which the qualitative character of $\mathcal{F}$ changes: below τ, the uninstantiated phase $\psi = 0$ is the unique minimum; above τ, it is a maximum and two new minima appear. The moment of change is the **bifurcation point**, the value of T at which $\psi = 0$ transitions from stable to unstable.

The characterisation of critical points as minima or maxima requires the second variation. In single-variable calculus, the second derivative test distinguishes a local minimum ($f'' > 0$) from a local maximum ($f'' < 0$) at a critical point. For a functional, the analogous test is the second functional derivative evaluated at the critical configuration. If $\delta^2\mathcal{F}/\delta\psi^2\big|_{\psi_0} > 0$, the functional is locally convex at ψ_0: the configuration is at a stable minimum. If negative, the functional is locally concave and the critical point is an unstable maximum.

The critical point of interest is $\psi = 0$, the uninstantiated configuration. To determine whether this is modally stable (the precausal phase is robust) or unstable (it cannot be maintained), one must examine the curvature of $\mathcal{F}$ at this point. Differentiating $\delta\mathcal{F}/\delta\psi = -2\varepsilon\psi + 4b\psi^3$ once more with respect to ψ and evaluating at $\psi = 0$ eliminates the cubic term entirely, since it contributes only $12b\psi^2\big|_{\psi=0} = 0$. What remains is purely the contribution of the quadratic term:

$$\left.\frac{\delta^2\mathcal{F}}{\delta\psi^2}\right|_{\psi=0} = -2\varepsilon(T) = 2(\tau - T) \tag{15}$$

This is positive for $T < \tau$ (i.e., $\varepsilon < 0$), zero for $T = \tau$, and negative for $T > \tau$ (i.e., $\varepsilon > 0$). The second variation vanishes at $T = \tau$, and this vanishing is the mathematical realisation of the threshold: the precise tension value at which the uninstantiated configuration ceases to be stable. What remains is to give τ a rigorous definition that does not presuppose its own value.

The naive approach would be to define τ as the solution to $\delta^2\mathcal{F}/\delta\psi^2\big|_{\psi=0} = 0$, which from the expression above gives $-2\varepsilon = 0$ and thus $T = \tau$, a circularity that conceals rather than reveals. The correct approach is to observe that τ is defined by a property of the functional, not by an algebraic equation: it is the value below which $\psi = 0$ is a stable minimum and above which it is not. This is precisely the infimum of the set of tension values for which the second variation vanishes:

$$\tau = \inf\left\{T(C) : \left.\frac{\delta^2\mathcal{F}}{\delta\psi^2}\right|_{\psi=0} = 0\right\} \tag{16}$$

The use of the infimum rather than a simple equation reflects the possibility that configurations at exactly $T = \tau$ form a set of measure zero in $\mathcal{C}$: the threshold is the greatest lower bound of all tension values that can trigger instantiation, and this bound may or may not be achieved by any individual configuration. In practice, when the functional takes the quadratic-quartic form above, the infimum is achieved and τ is well-defined as a specific value.

The significance of this derivation is that τ emerges from the structure of $\mathcal{F}$ itself. It is not a parameter that the theorist chooses. Once the functional is specified, and its form is constrained by the symmetries of the precausal substrate and the requirement that the transition be second order: the threshold is fixed. In this sense, the modal threshold is not a contingent feature of the theory. It is a structural necessity: any substrate governed by a modal potential of this form will have a threshold, and that threshold will be the bifurcation point of the functional.

The passage from $\psi = 0$ to the bifurcated minimum is modal sublimation, with no temporal unfolding. The next section characterises that transition in detail.

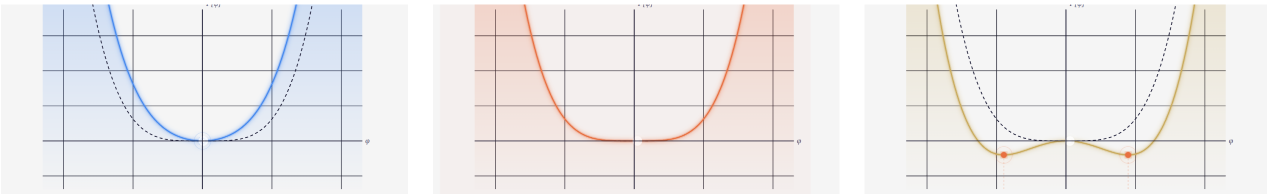

Figure 1: The modal potential $F[\varphi]$ at three values of excess tension ε. *Left:* $\varepsilon = +0.8$, single valley at $\varphi = 0$; the substrate configuration rests at the unique minimum, uninstantiated. *Centre:* $\varepsilon = 0$, the modal threshold; the valley flattens and the minimum is about to split. *Right:* $\varepsilon = -0.8$, two degenerate minima at $\varphi = \pm\sqrt{-\varepsilon}$; the configuration at $\varphi = 0$ is now a local maximum and cannot remain there. The dashed curve is the reference potential at $\varepsilon = 0$.

3.6. Layman's Summary

Not every tension imbalance produces a universe. This section derives a critical tension value, the modal threshold τ, below which the substrate remains abstract and unobservable, and above which it cannot remain that way. Think of water vapour: it can exist invisibly up to a certain temperature, but above that point it must condense. Below τ, the configuration has a single stable resting state with no geometry; above τ, that resting state becomes unstable and geometry must appear. The threshold is not assumed: it is calculated as the exact tension value at which the mathematical potential changes from single-welled to double-welled, and its existence is a structural necessity rather than a postulate.

4. Modal Sublimation

The transition induced by crossing the modal threshold is unlike any physical process encountered in the instantiated world. It has no duration, no mechanism, no intermediate states, and no energy cost. It is not that the transition is very fast; temporal concepts simply do not apply to it. The precausal substrate is a nonspatiotemporal domain: there is no time in which a transition could unfold, and no space across which it could propagate. What occurs at the threshold is not a process but a **modal reorganisation**, a shift in the mode of being of the configuration from uninstantiated to instantiated.

This transition is called **modal sublimation**, by analogy with the physical phenomenon in which a solid passes directly into a gaseous state, bypassing the liquid phase. In modal sublimation, every intermediate mode is bypassed because only two modes of being are available: uninstantiated existence within the precausal substrate, and instantiated existence as geometry. There is no intermediate mode that is partially spatial or partially causal.

Modal sublimation is characterised by four properties:

- **Nontemporal**: no before or after, no duration, no sequence of states; the configuration transitions *into* time rather than *through* it
- **Nonmechanistic**: no propagating influence, no force, no interaction
- **Nonenergetic**: no conservation law, no energy input or output

- **Strictly modal**: a change in the mode of being, not a change of state within an existing mode

Each of these four properties requires unpacking, because each contradicts a default assumption carried over from the instantiated world.

4.1. The nontemporal character

The nontemporal character is the most radical. In everyday physics, every process unfolds in time: there is a before and an after, a cause and an effect, a duration that can in principle be measured. Even processes described as instantaneous in non-relativistic mechanics, a collision, a decay, a measurement, are idealised limits of processes that have a finite, if arbitrarily small, duration. Modal sublimation is not a limit of this kind. It does not occur quickly; it does not occur in time at all. The configuration does not spend even an infinitesimal moment in transition. The reason is not dynamical but ontological: time is a structure of the instantiated domain, and the transition is the very event that brings that domain into being. To ask how long modal sublimation takes is to apply a concept to a context in which the concept has not yet been defined. This is not a failure of measurement precision; it is a logical impossibility of the same kind as asking what temperature a mathematical proof achieves. The configuration does not transition *through* time; it transitions *into* time, carrying time as a consequence of instantiation rather than as a pre-existing container in which the transition unfolds.

4.2. The nonmechanistic character

The nonmechanistic character follows directly from the nontemporal character. In the instantiated domain, every change of state is mediated by an interaction: a force, a field, a messenger particle, a propagating influence of some kind. These mediators all require spacetime in which to propagate and physical degrees of freedom to couple to. The precausal substrate has neither. No interaction can mediate the modal sublimation because interactions are structures of the domain that sublimation brings into being. There is no force that pushes a configuration across the threshold; the threshold is not a barrier that must be overcome but a logical boundary at which a particular mode of being ceases to be available. The transition is not driven by anything external. It is a logical consequence of the modal structure of the configuration itself, in the same way that the truth of a mathematical theorem is not caused by anything but follows necessarily from the axioms. This has a crucial implication: PST does not require a mechanism, an initial impetus, or a first cause. The demand for a mechanism is itself an instantiated-world concept, and applying it to the transition that precedes instantiation is a category error of the same kind as the causal question discussed in Section 5.

4.3. The nonenergetic character

The nonenergetic character removes the last temptation to treat modal sublimation as a physical event in disguise. Energy is a conserved quantity defined within the instantiated domain: it requires a Hamiltonian, a time-translation symmetry via Noether's theorem, and ultimately a spacetime manifold on which these structures can be defined. None of these are available in the precausal substrate. The relevant conservation statement in PST is not energetic but modal: it is the preservation of the asymmetric tension functional $T(C)$ across the threshold. The tension that characterises the precausal configuration is the same quantity that, after projection via Π, appears as the stress-energy tensor $T_{\mu\nu}$ in the instantiated domain, as shown in Section 7. The Einstein field equations, on this reading, are not equations of motion governing a physical process; they are a **conservation statement** expressing that the structural content of the precausal configuration is faithfully preserved in the instantiated geometry. Nothing is created across the threshold; nothing is destroyed. The same structural content is expressed in a different modal register, the way that the same logical argument can be expressed in different formal languages without gaining or losing content.

4.4. The strictly modal character

The strictly modal character is the most philosophically precise formulation of what the other three properties imply. A change of state *within* an existing mode, a particle moving, a field oscillating, a temperature rising, leaves the mode of being intact while altering something within it. The ontological register is unchanged: the system is still a physical system in spacetime, governed by the same causal laws, measured in the same units. Modal sublimation changes the mode itself. The configuration does not acquire new properties within the precausal substrate; it changes the ontological register in which all its properties are expressed. Before sublimation, the configuration's structural content exists as modal tension in a nonspatiotemporal domain. After sublimation, that same structural content exists as geometric curvature and matter-energy in a Lorentzian manifold. There is no moment at which the configuration is partially in each mode; the transition is all-at-once in the modal sense, even though "all-at-once" is itself a temporal phrase that strictly applies only after time has come into being. This is the precise sense in which PST's fundamental transition is irreducible to any physical process: physical processes are changes of state within the instantiated mode, and modal sublimation is the event that instantiates the mode itself.

The result of this transition is a **geometry**: a Lorentzian manifold (M, g) that becomes the spacetime background for all subsequent physical processes. Geometry does not emerge *within* spacetime; it *becomes* spacetime in the act of sublimation.

4.5. Categorical formulation

The mathematical language suited to a transition between fundamentally different types of structure is category theory, specifically the theory of functors between categories [9]. Construct two categories as follows. The **precausal category S** has as objects all configurations $C \in \mathcal{C}$ with $T(C) \geq \tau$, and as morphisms distinction preserving maps $f : C \to C'$. The **geometric category G** has as objects Lorentzian manifolds (M, g) and as morphisms smooth maps preserving the causal structure.

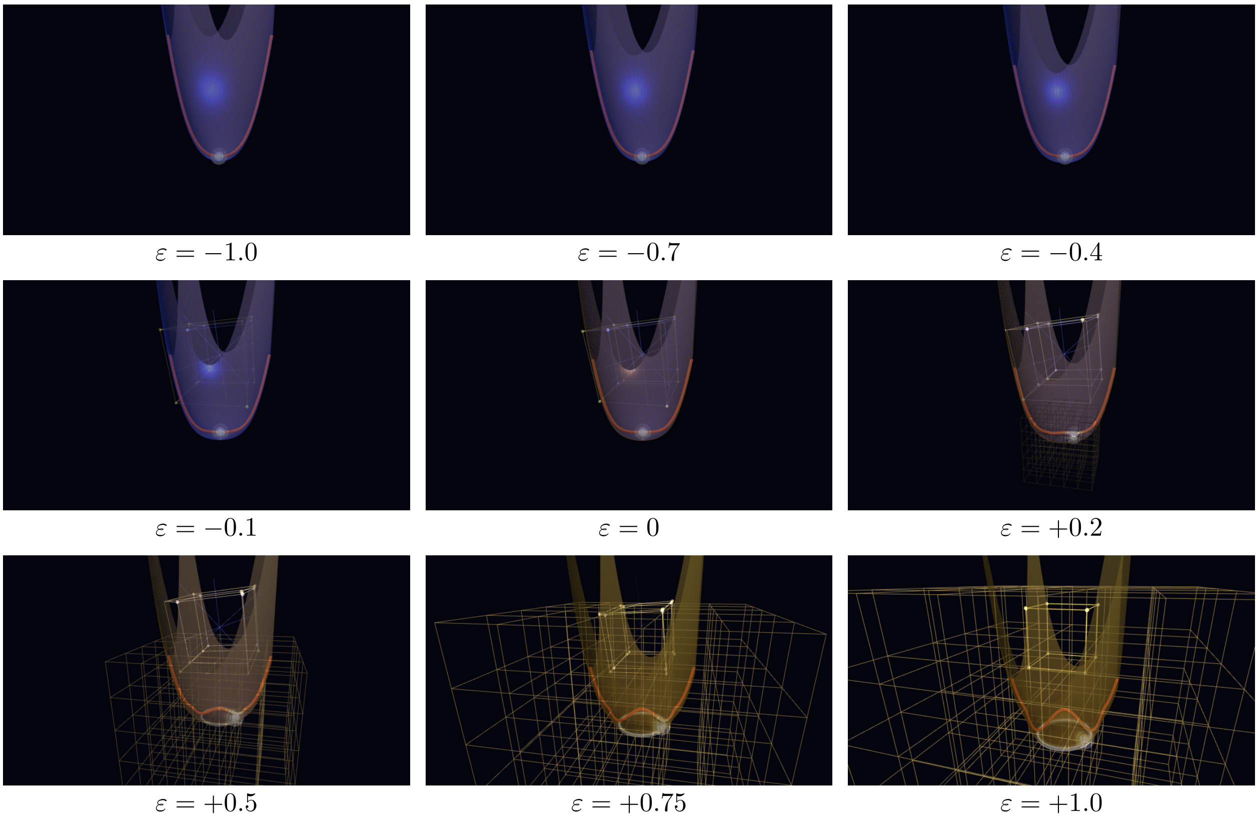

Figure 2: Visual model of modal sublimation across nine values of the excess tension $\varepsilon = T(C) - \tau$ (Section 3): $\varepsilon < 0$ precausal, $\varepsilon > 0$ instantiated. Each frame shows the modal potential $F[\psi, \varepsilon] = -\varepsilon\psi^2 + \frac{1}{4}\psi^4$ rendered as a rotationally symmetric surface (lower portion), with the orange tube tracing its one-dimensional cross-section and the white sphere marking the position of the order parameter ψ at the current potential minimum. *Top row* ($\varepsilon < 0$): the precausal substrate; the potential forms a deep bowl with a single minimum at $\psi = 0$, the order parameter is confined to the origin, and no spacetime grid is present. *Middle row* ($\varepsilon \approx 0$): approaching and crossing the modal threshold; the bowl flattens and its centre begins to rise, the tesseract (4-cube) becomes visible as it begins its dimensional collapse, and a nascent spacetime lattice emerges. *Bottom row* ($\varepsilon > 0$): instantiated geometry; the potential has acquired a continuous ring of degenerate minima (the vacuum manifold $\mathcal{V} \cong S^1$) visible at the base of the Mexican-hat surface, the order parameter has settled at one point on this ring, the tesseract has collapsed to a 3-cube, and the golden spacetime grid fills the surrounding space.

Modal sublimation is the functor:

$$\Phi : \mathbf{S} \to \mathbf{G} \tag{17}$$

with four properties encoding the fundamental features of the transition:

1. Φ **exists at the categorical boundary**; it is not a morphism within $\mathbf{S}$ or $\mathbf{G}$, formalizing the nonmechanistic character of the transition
2. Φ **preserves tension as curvature**: $T(C) \mapsto \kappa(\Phi(C))$, where κ is the scalar curvature of the instantiated manifold
3. Φ **is surjective**: every Lorentzian manifold arises as the image of some precausal configuration
4. Φ **has no inverse**: once a configuration is instantiated, the geometric domain has no access to the precausal substrate that produced it

4.6. Explicit construction of Φ

The functor is constructed in three steps.

Step 1, Geometric realization. The manifold M and the map ρ are not given independently before the construction; they are jointly determined by it.

A legitimate concern requires attention before the construction proceeds. The notation $\rho : D \to M$ appears to presuppose M, a set of spacetime points, before the construction that is supposed to *produce* spacetime has been carried out. If spacetime points are already available as the target of ρ, then the theory has not derived spacetime from the substrate; it has merely labeled pre-existing points with properties from D. This would be a foundational circularity.

The resolution is that M is not given prior to and independently of ρ. The construction proceeds in the following order. First, note that any surjective function $\rho : D \twoheadrightarrow X$, for *any* set X, partitions D into equivalence classes: $a \sim_\rho b$ if and only if $\rho(a) = \rho(b)$. The equivalence classes $[a]_\rho = \{b \in D : \rho(b) = \rho(a)\}$ are the fibres of ρ, and their collection $D/{\sim_\rho}$ is a quotient set that depends only on D and the partition, not on any property of X. A spacetime *point* is then defined to be a fibre of ρ: it is an equivalence class of properties, not a pre-existing geometric entity. The manifold M is the quotient set $D/{\sim_\rho}$ equipped with whatever additional structure (topology, differentiable structure, metric) the subsequent steps impose. The map ρ is thus a *quotient map*, and M is a *derived* set: it comes into existence as the image of ρ, rather than being a container that ρ maps D into.

This construction has a philosophical precedent in Leibniz's principle of the identity of indiscernibles: things that share all the same properties are the same thing. Here, the converse is operative, properties that are assigned to the same geometric locus by ρ are identified, and the locus is nothing over and above that identification. A spacetime point has no content beyond the property-bundle it collects; it is not an additional ontological entity.

More concretely: M is characterised as *any* smooth manifold of dimension n whose topology and differential structure are consistent, a posteriori, with requiring that the metric constructed in Step 2 is well-defined and smooth. Among all candidate manifolds satisfying this consistency condition, the tension structure of (D, δ) further constrains which are admissible: a candidate M is admissible if and only if there exists a partition of D whose induced tension structure $T(\{[a], [b]\}) = T(\{a, b\})$ reproduces the relational data already present in (D, δ). This constraint, not an external stipulation, selects the spacetime topology and dimension.

Given a choice of admissible M, define the **realization map** $\rho : D \to M$ as the quotient projection sending each property $a \in D$ to its equivalence class $\rho(a) = [a]_\rho \in M$. The nonuniqueness of ρ, the freedom to partition D into fibres in multiple ways that are all consistent with the tension structure, is precisely the diffeomorphism invariance of General Relativity [1, 10]. It emerges here as a consequence of the nonuniqueness of geometric realization from the pre-geometric substrate, rather than as a separately imposed symmetry. There is no canonical labeling of spacetime points: this is not a deficiency of the theory but a structural result, matching the gauge freedom of GR.

Step 2, Tension-to-metric map. Introduce the **null threshold** $\tau_{\text{null}} \in \mathbb{R}^+$ as a second structural input alongside τ. While τ is a condition on the global tension $T(C)$, τ_{null} is a condition on pairwise tensions $T(\{a, b\})$: it is the value at which the two-property tension is null, i.e., at which $g = 0$ by definition. The two thresholds operate at different levels and their relationship is an open problem listed in Section 14. The metric is then:

$$g_{\mu\nu}(\rho(a), \rho(b)) = \eta_{\mu\nu} \cdot \text{sgn}(T(\{a, b\}) - \tau_{\text{null}}) \cdot |T(\{a, b\}) - \tau_{\text{null}}|^{1/2} \tag{18}$$

where $\eta_{\mu\nu}$ is a placeholder for the diagonal sign pattern $\text{diag}(-1, +1, +1, +1)$ that the construction induces, not a presupposition of Minkowski geometry. Its appearance here is definitional: the signature theorem below shows that the sign pattern is fixed by the tension structure; $\eta_{\mu\nu}$ names that pattern. This gives:

- $T(\{a, b\}) = \tau_{\text{null}} \;\Rightarrow\; g = 0$: null separation
- $T(\{a, b\}) > \tau_{\text{null}} \;\Rightarrow\; g > 0$: spacelike
- $T(\{a, b\}) < \tau_{\text{null}} \;\Rightarrow\; g < 0$: timelike

The Lorentzian signature is not postulated; it follows from the bifurcation of tension values around τ_{null}.

Two clarifications are necessary. First, equation (18) is strictly a pairwise assignment: it attaches a signed scalar to each pair of realised points, not a $(0, 2)$-tensor to a tangent space. The metric tensor field $g_{\mu\nu}(x)$ is recovered by the continuum limit: take a and b to be nearby properties such that $\rho(a)$ and $\rho(b)$ approach x along a tangent vector v^μ; then $g_{\mu\nu}(x)\, v^\mu v^\nu$ arises as the limit of equation (18) along the chord, analogous to recovering a Riemannian metric as the Hessian of a distance function. Second, the magnitude $\left|T(\{a, b\}) - \tau_{\text{null}}\right|^{1/2}$ is the minimal-regularity choice: a linear dependence would give a metric with a non-smooth gradient at the null cone, while the square root yields Hölder-$\frac{1}{2}$ behaviour there, the appropriate regularity class for a structure that degenerates at $T = \tau_{\text{null}}$ and which ensures that equation (51) blows up precisely at the modal null threshold rather than at $T = 0$.

Step 3, Action on morphisms. For $f : C \to C'$ in $\mathbf{S}$, $\Phi(f) : (M, g) \to (M', g')$ is the diffeomorphism induced by f via ρ. Distinction preserving maps in $\mathbf{S}$ correspond precisely to causal structure preserving maps in $\mathbf{G}$, ensuring functoriality.

The projection operator. Define the **tension projection operator** Π:

$$\Pi(T(C))(x) = \int_{\mathcal{C}} T(C') \cdot K(x, C')\, d\mu(C') \tag{19}$$

where $K(x, C') = \delta(x - \rho(C'))$ for local configurations. The role of Π here is structurally analogous to the Dirac operator in Connes' noncommutative geometry [49]: in both cases a single operator encodes the geometric and algebraic content of the instantiated domain from an underlying non-spatial structure. The stress-energy tensor at $x \in M$ in the isotropic, homogeneous case is:

$$T^{(\text{iso})}_{\mu\nu}(x) = \Pi(T(C))(x) \cdot g_{\mu\nu}(x) \tag{20}$$

This isotropic scalar form $T^{(\text{iso})}_{\mu\nu} = \Pi(T(C))\, g_{\mu\nu}$ gives a stress-energy tensor proportional to the metric, which is the correct structure for a cosmological constant ($T_{\mu\nu} = \Lambda g_{\mu\nu}$) but not for general matter distributions. Pressurised matter, radiation, and anisotropic fields all have $T_{\mu\nu} \neq \lambda g_{\mu\nu}$ for any scalar λ. The general form requires Π to be a tensor-valued projection $T_{\mu\nu}(x) = \Pi_{\mu\nu}(T(C))(x)$, whose components encode the directional distribution of tension in the precausal configuration. The scalar form is the leading-order approximation valid for isotropic configurations; the full tensor generalisation is an open problem in Section 14. The noninvertibility of Φ is immediate: ρ discards the internal structure of D that does not affect pairwise tensions. Geometry encodes the projection of T through Π but not the full precausal configuration that generated it, just as a shadow records an object's outline but not its internal structure.

4.7. Uniqueness of Φ and the origin of diffeomorphism invariance

The construction proceeds via an explicit choice of realization map $\rho : D \to M$, which is not determined by the precausal structure alone. This raises the question of how strongly the output functor Φ depends on the choice of ρ, and what that dependence implies for the physical content of the theory. The answer is a uniqueness theorem: Φ is unique up to diffeomorphism, and the ambiguity in Φ is completely and exactly characterised by diffeomorphism invariance, with no residual freedom beyond it.

The proof follows directly from the construction. Given ρ, the metric $g_{\mu\nu}$ at every pair of points is fully determined by equation (18): there are no further choices once the realization map is fixed. Now suppose $\rho' = f \circ \rho$ for some smooth bijection $f : M \to M'$. The tension values $T(\{a, b\})$ depend only on $a, b \in D$, not on their images under ρ, so they are unchanged by the substitution. Therefore the metric computed at $(\rho'(a), \rho'(b))$ via equation (18) satisfies

$g'(\rho'(a), \rho'(b)) = g(\rho(a), \rho(b))$, which is exactly the condition that (M', g') is the pushforward of (M, g) along f. Any two choices of ρ therefore produce Lorentzian manifolds that are diffeomorphic to one another, together with a corresponding natural isomorphism between the two functors they define. Φ is unique up to this natural isomorphism, which at the level of manifolds is a diffeomorphism.

The physical significance of this result is that diffeomorphism invariance in General Relativity is not an additional symmetry requirement that must be imposed upon the theory by hand. It is the complete characterisation of the non-uniqueness of Φ: the relabelling freedom in choosing ρ survives into the instantiated manifold as coordinate freedom, and coordinate freedom is diffeomorphism invariance. There is no finer ambiguity and no coarser one; the two are identical. The question raised in the open problems section, whether Φ is unique and what physical symmetries characterise its ambiguities, is therefore answered here: Φ is unique up to diffeomorphism, its ambiguity is completely characterised, and the physical symmetry corresponding to that ambiguity is diffeomorphism invariance.

A complementary uniqueness result concerns the metric signature. The argument in Step 2 assigned the Minkowski signature $(-, +, +, +)$ through equation (18), on the grounds that asymmetric tension bifurcates about a null threshold τ_{null} separating spacelike from timelike separations. But one may ask whether the signature, and in particular the number of negative eigenvalues, is fixed by the precausal structure or whether it could vary across different instantiation events. The answer is that it is fixed uniquely. A precausal configuration C is asymmetric by the definition of property differentiation: $T(C) \neq T(\bar{C})$, because asymmetric tension distinguishes a configuration from its complement. This asymmetry means the tension spectrum is not symmetric about τ_{null}, and consequently there is precisely one direction in configuration space along which the pairwise tension is minimal and crosses τ_{null} from below. It is this direction that maps to the temporal dimension of the instantiated manifold. All other independent directions carry pairwise tensions above τ_{null} and therefore map to spacelike dimensions. The signature $(-, +, +, +)$ is uniquely fixed because asymmetric tension has exactly one extremal direction. Different choices of ρ may permute the spatial coordinates or reverse their orientations, but they cannot change the count of temporal dimensions, which is fixed by the tension structure of (D, δ, T) itself. The Lorentzian signature is not postulated; it is derived, and it could not have been otherwise.

4.8. Layman's Summary

Modal sublimation is the moment the threshold is crossed, not a moment in time, but a logical transition. When a configuration's tension exceeds τ, it cannot remain in its abstract precausal state; it is compelled, by the same kind of necessity that makes a mathematical theorem true, to become concrete. The result is a four-dimensional spacetime with the Lorentzian signature that physics recognises. This section constructs the map from the precausal domain to the geometric domain, shows it is unique, and proves that the coordinate-freedom of General Relativity (diffeomorphism invariance) is a consequence of that uniqueness rather than a separate postulate.

5. Coinstantiation of Spacetime and Causality

The previous section established that modal sublimation is a nontemporal, nonmechanistic transition: the functor $\Phi : \mathbf{S} \to \mathbf{G}$ maps a precausal configuration directly into a Lorentzian manifold without passing through any intermediate state and without requiring a prior temporal structure in which the transition could unfold. A consequence of this that deserves its own careful treatment is the status of causality itself. Causality is not, in PST, a background structure that was already present when the universe formed and that can then be applied to the event of its formation. Causality is **coinstantiated** with geometry: it arises at the same moment as the Lorentzian manifold that carries it, because it is constitutively defined in terms of that manifold. There is no moment, however brief or abstract, at which a causal relation exists but spacetime does not. The two are ontologically inseparable.

This point carries weight precisely because it is not obvious. Much of the philosophical and cosmological literature treats causality as prior to or more fundamental than spacetime. In the Kantian tradition, causality is a category of the understanding applied to experience, logically prior to any empirical fact about the world [35]. In causal set theory, the discrete partial order is the primary structure from which the continuum metric is to be recovered [4, 36]. In quantum gravity approaches that proceed via sum over histories, the causal structure of each history is taken as given, and the amplitude is computed over those histories [37]. Even in ordinary language, the question "what caused X?" is considered answerable in principle for any event X, the presumption being that every event has a cause and that the causal relation is always available to be invoked.

PST disagrees at the foundational level. The precausal substrate, as defined in Section 2, does not possess a causal order. It possesses property differentiation and the tension functional. These are not impoverished substitutes for causality; they are genuinely different kinds of structure. The distinction relation δ is symmetric and irreflexive: it registers that two properties differ from each other without specifying which is prior to the other or whether either depends on the other. The tension functional T assigns a magnitude to a configuration without implying that one configuration succeeds or precedes another. Neither δ nor T contains the ingredients from which a temporal order could be constructed. A temporal order requires at minimum an asymmetric, transitive relation on a set of events. The precausal substrate has none. Causality is not absent from the substrate in the way that a feature might be absent from a description that was meant to include it; it is absent from the substrate in the sense that the substrate belongs to a logical type to which causal relations do not apply.

5.1. Mathematical formulation of coinstantiation

Let $(M, g) = \Phi(C)$ be the Lorentzian manifold produced by modal sublimation of a configuration $C \in \mathbf{S}$. The causal structure on (M, g) is the partial order $\leq$ defined pointwise by the lightcone

v9.9

structure of g:

$$p \leq q \iff q \in J^+(p, g) \tag{21}$$

where $J^+(p, g)$ is the causal future of p under the metric g:

$$J^+(p, g) = \{\, q \in M \;\mid\; \exists \text{ future-directed causal curve } \gamma : [0,1] \to M,\ \gamma(0) = p,\ \gamma(1) = q \,\} \tag{22}$$

A curve is causal under g if and only if its tangent vector $\dot{\gamma}$ is everywhere non-spacelike, that is $g(\dot{\gamma}, \dot{\gamma}) \leq 0$ at every point. This definition depends on g in an essential way: the notion of a causal curve, and hence of causal precedence, is meaningless without a metric. Change the metric and the causal structure changes with it; remove the metric and the causal structure does not become unspecified but ceases to exist as a category of object.

The partial order $\leq$ in equation (21) has no pre-image in the precausal category **S**. Formally, there is no morphism or relation in **S** that maps to $\leq$ under Φ, because the objects of **S** are configurations equipped only with δ and T. This is not a deficiency of the model; it is a structural theorem: the functor Φ is not surjective on structure but on objects. It produces a manifold but does not carry into that manifold a pre-existing causal order; the causal order is generated anew by the Lorentzian signature of the metric produced by equation (18). Therefore the following holds:

$$\text{causality} \;=\; f(g) \;=\; f(\Phi(C)) \tag{23}$$

Causality is a derived relation, defined in terms of the metric, which is itself defined in terms of the precausal tension structure via Φ. It is not an independent primitive.

5.2. The direction of time as a structural consequence

The Lorentzian signature of the metric $(-, +, +, +)$ introduces an asymmetry that has no parallel in a purely Riemannian geometry. In a Riemannian manifold, all directions from a point are spacelike and the notion of a preferred future does not arise. In a Lorentzian manifold, the lightcone at each point divides the tangent space into two timelike halves, one conventionally designated as the future and the other as the past, and a continuous causal vector field on M selects one half globally, providing a **time orientation** [38, 39].

In PST this asymmetry is not imposed as an additional datum. It is inherited from the asymmetry of the precausal substrate. Asymmetric tension, by definition, is a non-symmetric functional: $T(C) \neq T(\bar{C})$ for all configurations C, where $\bar{C}$ is the complement of C in D. This asymmetry is preserved across the functor Φ as the Lorentzian signature, and the preferred time direction of the manifold reflects the direction of the residual tension gradient of the precausal configuration. The thermodynamic arrow of time, the cosmological arrow of time, and the distinction between retarded and advanced solutions in electrodynamics are all grounded in this single asymmetry inherited from the substrate.

A technical caveat is required. The complement-asymmetry $T(C) \neq T(\bar{C})$ is a pairwise condition; a global time orientation on M requires a continuous nowhere-vanishing timelike vector field on the entire manifold. By Geroch's theorem [38], a Lorentzian manifold is time-orientable if and only if such a field exists, which is a global topological condition not automatic from local asymmetry. In PST the global character of the precausal configuration C provides the resource: the tension gradient $\nabla_{\mathcal{C}} T$ evaluated in the preferred direction identified in Section 4 gives a precausal structure that Φ maps to a candidate timelike field on M. That this field is continuous and nowhere-vanishing, i.e. that global time-orientability follows from the precausal structure without additional assumption, is an open problem listed in Section 14.

5.3. The causal order and the degenerate vacuum

A further consequence of coinstantiation concerns the vacuum manifold. As established in Section 3, the post-threshold vacuum $\mathcal{V} \cong S^1$ is a continuum of configurations that minimize the modal potential at equal tension. Motion along $\mathcal{V}$ is the only stable state the vacuum admits. In the instantiated geometry, this motion projects under Π into persistent non-vanishing stress-energy, which in turn sustains curvature. Crucially, the causal order $\leq$ treats this motion as part of the background structure: the lightcone structure generated by the projected metric encodes the vacuum motion as a baseline curvature rather than as a deviation from flatness. This is why orbital motion at all scales, from subatomic to galactic, does not require an initial condition that specifies it separately. It is the ground state of the causal structure itself, not a perturbation applied to an otherwise static manifold.

5.4. Category error and the dissolution of first-cause questions

With coinstantiation established, we are in a position to address what is perhaps the most persistent question in the philosophy of cosmology: what caused the universe to come into being? The question appears legitimate. It has the grammatical form of a causal question, it points to a definite event (the coming-into-being of the universe), and it invites an answer of the form "X caused the universe".

The question is not, however, answerable. It is not unanswerable in the sense that the answer is hidden from us or that the evidence has been destroyed; it is unanswerable because it commits a **category error**, applying a relation to an entity that belongs to a logical type to which the relation does not apply [40, 41]. To ask what caused the universe is to presuppose that the causal relation exists and is applicable to the event of the universe's formation. But the causal relation, as shown above, is coinstantiated with the universe: it comes into being with the very event whose cause is being sought. A relation cannot serve as the explanation for the event that constitutes its own first instantiation.

The parallel with simpler category errors is instructive. "What is north of the North Pole?" presupposes that the directional relation north-of is defined at the North Pole, but the relation is defined only within the coordinate system that has the North Pole as its boundary. "What was there before the Big Bang?" presupposes that the relation earlier-than is defined at the Big Bang, but that relation depends on the spacetime whose earliest moment is the Big Bang. "What caused the universe?" presupposes that the causal relation is defined for the event of universal instantiation, but that relation depends on the very spacetime that is being instantiated. In each case the error is structural: the question applies a relation in a domain where that relation does not exist.

The dissolution of the category error is not a retreat. PST does not avoid the question by refusing to engage; it dissolves it by exhibiting the precise logical structure that makes it malformed. This dissolution is, moreover, constructive: in place of the unanswerable question, PST offers the answerable one. Why does the universe exist at all? Because the precausal substrate possesses property differentiation, and property differentiation entails asymmetric tension, and once tension exceeds the modal threshold, modal sublimation to instantiated geometry is logically necessary, not contingently triggered. The question Leibniz posed, "why is there something rather than nothing?" [42], receives not an arbitrary answer but a principled one: because the logical possibility of property differentiation is irreducible, and a substrate bearing irreducible property differentiation cannot remain uninstantiated once the threshold condition is satisfied. Existence is not a brute fact; it is the necessary expression of a logically unavoidable modal imbalance.

5.5. Contrast with competing approaches

Several approaches in the literature address the question of cosmological origins in ways that are instructive to compare with PST.

The **Hartle-Hawking no-boundary proposal** [37] avoids a boundary in time by making the metric Euclidean in the deep past, eliminating the initial singularity by changing the signature. This is technically elegant but it does not dissolve the category error: the sum over geometries is still computed within a pre-existing mathematical framework equipped with a notion of quantum amplitude, and the question of why that framework is instantiated rather than another is not addressed. The no-boundary proposal also treats causality as part of the background formalism rather than as a derived structure.

Vilenkin's tunneling proposal [43] posits a quantum tunneling event from nothing. But "nothing" in this context is not the absence of structure; it is a specific quantum gravitational state with a precisely defined wave function. The transition from that state presupposes both quantum mechanics and the causal structure within which the transition is defined. The proposal is internally consistent but does not escape the regress: it pushes the question back to the prior state and the framework defining the transition.

Causal set theory [4] makes the causal partial order primitive and attempts to derive the continuum from it. This is in some respects the inverse of PST: where causal set theory starts with $\leq$ and tries to recover g, PST derives both g and $\leq$ from Φ. Causal set theory faces the question of what determines the particular causal set that becomes our universe, a question it answers by probabilistic sampling over a measure on causal sets. PST does not face this question in the same form: the particular geometry that arises is determined by the tension configuration C, which is itself subject to the modal potential.

The position of PST is that all approaches that take causality as a starting point, whether as a background relation, a quantum mechanical framework, or a primitive partial order, are working within the instantiated domain and can at best describe the structure of what has been instantiated. The question of why anything is instantiated at all requires a theory that can speak from outside the instantiated domain. The precausal substrate, and the functor Φ through which it gives rise to geometry, is that theory.

5.6. Layman's Summary

Spacetime and the causal order (the fact that causes precede effects) do not arise one after the other. They are two aspects of a single transition and cannot exist independently of each other. This has a profound consequence: there is no "before" the universe, because time itself is part of what the transition produces. Asking what caused the universe is a category error, like asking what is north of the North Pole. The Big Bang is not a beginning at a point in time; it is what the transition looks like from inside the geometry it creates.

6. The Emergence Chain

The preceding four sections have built the foundations of PST step by step: from the single primitive of property differentiation, through asymmetric tension and the modal threshold, to modal sublimation and the coinstantiation of geometry and causality. This section draws those steps together into a single, unified chain of logical entailment and examines what that chain is, what kind of necessity binds each step to the next, what distinguishes it from competing accounts of emergence in physics, and what it means for the chain to be complete.

The table below summarises the eight steps that constitute the chain. Each entry names the emergent concept, gives the mathematical object that represents it, and implies a link to its predecessor. The prose that follows unpacks each link in turn.

Step	Concept	Mathematical object
1	Property differentiation	(D, δ), a set with a symmetric, irreflexive distinction relation
2	Asymmetric tension	$T : \mathcal{C} \to \mathbb{R}^+,\ T(C) \neq T(\bar{C})$ for all C
3	Modal threshold	Bifurcation point τ of $\mathcal{F}(\psi, C)$, eq. (16)
4	Modal sublimation	Functor $\Phi : \mathbf{S} \to \mathbf{G}$, constructed via ρ and eq. (18)
5	Geometry as spacetime	Lorentzian manifold (M, g) with signature derived from τ_{null}
6	Causality	Partial order $\leq$ on M via $J^+(p, g)$; no pre-image in $\mathbf{S}$
7	Energy, matter, fields	$T_{\mu\nu} = \Pi(T(C))$; $G_{\mu\nu} = 8\pi G \cdot \Pi(T(C))$; bosonic field sector (fermionic matter contingent on the spin-statistics open problem)
8	Angular momentum	Noether charge $L = 2r_0^2\,\dot{\theta}$ of the U(1) symmetry of the vacuum manifold $\mathcal{V} \cong S^1$

6.1. The character of the chain

The first thing to establish is what kind of chain this is. It is not a causal chain. The steps do not follow each other in time; they are not events that happen sequentially in a universe that already exists. It is not an evolutionary chain in the sense used in cosmology, where one physical state develops into another under the action of laws that are themselves already given. The chain is a chain of **logical entailment**: each concept is not merely correlated with its successor but is the necessary and sufficient condition for it. Given step n, step $n + 1$ is not optional; it is logically unavoidable. Remove any link and the entire structure downstream collapses.

This is a stronger claim than is usually made in emergence theories. In most accounts, emergence is a relation between levels of description: a higher-level description emerges from a lower-level one when the higher-level regularities can, in principle, be derived from the lower-level equations. The chain has a direction but not a necessity. In PST the chain has both. The direction runs from the logically irreducible toward the structurally complex, and at every step the direction is the only one available. The question "why does tension arise from differentiation?" does not have an empirical answer, the sort of answer one would give by pointing to a law of nature or a measured constant. It has a logical answer: any configuration of two or more distinct properties contains non-identity internal relations, and those relations cannot collectively neutralise. Tension is not observed to follow from differentiation; it is entailed by it, together with the complement-asymmetry condition.

A precise statement of what the chain entails requires distinguishing two levels. At the *structural* *level*, the steps are necessary given the primitive and the structural commitments of PST: tension follows from distinctness together with the complement-asymmetry condition; the threshold follows from the existence of tension together with the commitment to a bifurcation structure; modal sublimation follows from the threshold by the instability argument; spacetime and causality coinstantiate via the functor Φ. At the *architectural level*, the specific mathematical objects representing each step, the Landau-Ginzburg functional, the gradient flow, the canonical measure, are the minimal motivated choices consistent with the structural requirements; they are not strictly entailed by (D, δ) alone, as the epistemological caveat of Section 2 makes explicit. Steps 3 through 8 are necessary given the structural encoding, which is itself logically motivated rather than derived. The chain is closed at the structural level; it is architecturally motivated at the representational level. This two-tier character is not a weakness; it identifies precisely where the foundational programme remains open and where future derivations could replace motivated choices with strict entailments.

6.2. Step 1 to Step 2: from differentiation to tension

Property differentiation is the capacity for distinctions to exist: the set D equipped with the distinction relation δ, symmetric and irreflexive, requiring no loci, relata, or temporal structure. This is the single primitive from which the entire theory is built. Its logical irreducibility was established in Section 2: any attempt to explain why distinctions can exist already presupposes a context in which two things differ from each other, committing a circularity. The primitive cannot be grounded further without circularity; it can only be identified.

Asymmetric tension follows from differentiation together with the complement-asymmetry condition. A configuration $C \subset D$ drawn from a set of distinct properties admits internal non-identity relations. Since δ is symmetric, $\delta(a, b)$ and $\delta(b, a)$ hold simultaneously; these relations carry no preferred direction. The asymmetry in asymmetric tension is not directional but complementary: the structural condition that no configuration and its complement carry equal tension, $T(C) \neq T(\bar{C})$. This condition is the precise additional commitment that converts step 1 into step 2. It cannot be derived from (D, δ) alone; it is the assertion that the precausal substrate is generically asymmetric, and it functions as the bridge between the bare existence of distinctions and their structural expression as tension. Granting it, the move from step 1 to step 2 is structurally closed: step 2 is step 1 made structural, with the complement-asymmetry condition as the content of that structuring move. This is the foundational hinge of the chain.

6.3. Step 2 to Step 3: from tension to threshold

Asymmetric tension is present in every configuration of more than one distinct property. That it exists is already established by step 2. What step 3 adds is the recognition that tension admits a *critical value*, a point at which the structural consequences of tension change qualitatively rather than quantitatively. Below this critical value, the configuration can sustain a unique, stable modal position at $\psi = 0$: uninstantiated, without geometry, without causality. Above it, the configuration cannot. The threshold τ is not introduced as an external parameter but is derived as the infimum of the set of tension values at which the second variation of the modal potential functional $\mathcal{F}(\psi, C)$ vanishes at $\psi = 0$, equation (16).

The necessity of step 3 is that there must be such a threshold. The modal potential functional contains a quadratic term whose coefficient is proportional to $T(C) - \tau$, and a quartic term that bounds the functional below. The qualitative bifurcation between the single-minimum and double-minimum regimes is a structural consequence of the polynomial structure of $\mathcal{F}$, which in turn reflects the two competing tendencies present in any configuration: the tendency toward modal coherence (captured by the quadratic restoring term) and the tendency toward structural saturation (captured by the quartic term). These two tendencies are not postulated independently; they are the formal expression of the two irreducible aspects of any tension-carrying configuration. A threshold must exist because the quadratic and quartic terms cannot produce a bifurcation without a critical value at which the quadratic coefficient changes sign. That critical value is τ.

6.4. Step 3 to Step 4: from threshold to sublimation

The threshold establishes a condition; modal sublimation is what happens when that condition is met. Once $T(C) > \tau$, the configuration at $\psi = 0$ is a local maximum, not a minimum. The configuration cannot remain there. This is not a dynamical statement; there is no time over which the configuration moves. It is a logical statement: a configuration whose tension exceeds the threshold and whose order parameter is at the unstable critical point is internally incoherent. Its degree of instantiation is inconsistent with its tension. That incoherence is resolved not by a process but by a logical transition to the nearest coherent state.

That transition is the functor $\Phi : \mathbf{S} \to \mathbf{G}$, the map from the precausal category to the geometric category. The functor is not postulated; its existence follows from the threshold condition together with the definition of the two categories. The precausal category $\mathbf{S}$ has as objects configurations with tension above τ; the geometric category $\mathbf{G}$ has as objects Lorentzian manifolds. The functor Φ assigns to each configuration its geometrically realised image, constructed explicitly via the realization map ρ and the tension-to-metric assignment of equation (18). The move from step 3 to step 4 is thus the move from a necessary condition to its necessary consequence: if the threshold is the condition under which uninstantiated existence becomes logically incoherent, then modal sublimation is the resolution of that incoherence. Step 4 is not something that happens after step 3; it is what step 3 is, from the perspective of its consequence.

6.5. Step 4 to Steps 5 and 6: from sublimation to spacetime and causality

The image of Φ is a Lorentzian manifold (M, g). The Lorentzian signature is not assumed but derived: it follows from the bifurcation of pairwise tension values around the null threshold τ_{null}, as shown in equation (18). Pairs with tension above τ_{null} map to spacelike separations; pairs below to timelike separations; pairs at exactly τ_{null} to null separations. The $(-, +, +, +)$ signature of spacetime is therefore a direct expression of the way tension distributes across pairwise configurations in the precausal domain.

Causality arises simultaneously with the manifold. The causal order $\leq$ defined by $J^+(p, g)$ is constitutively tied to g: it cannot exist before or independently of the metric. Steps 5 and 6 are therefore not two sequential events but two aspects of a single transition. The reason they are listed separately in the table is conceptual, not temporal: each warrants individual attention because each dissolves a traditional presupposition. Step 5 dissolves the presupposition that geometry is given; step 6 dissolves the presupposition that causality is given. Both are derived, both from the same functor, and both are coinstantiated.

6.6. Steps 6 and 7: from causality and geometry to matter

Matter, energy, and fields are sometimes treated in foundational physics as the *content* of spacetime, distinct from the geometric container. PST dissolves this distinction. The stress-energy tensor $T_{\mu\nu}$ at any point $x \in M$ is the projection of the precausal tension functional through the operator Π, as in equation (24). It is not a separate input into the Einstein field equations; it is the same tension structure that produced the geometry, now expressed in geometric units via the projection. The field equations $G_{\mu\nu} = 8\pi G \cdot \Pi(T(C))$ are a conservation statement: they express the fact that the total tension content of the precausal configuration is faithfully preserved in the instantiated manifold, distributed as both curvature and stress-energy. Gravity and matter are not two things; they are two aspects of one projection.

The factor $8\pi G$ is not a fundamental constant of the precausal domain. It is a unit conversion factor introduced by Π, translating between dimensionless modal tension and physical stress-energy units. Its numerical value is a property of the projection map, not of the substrate. This reframes the longstanding puzzle of the gravitational constant from a question about nature to a question about the structure of the projection operator, a question that in principle has a derivable answer once the full form of Π is specified.

"Matter" at Step 7 currently means stress-energy and the bosonic field sector derived in Section 7. The fermionic sector, which constitutes the matter content of the Standard Model, requires half-integer spin and the spin–statistics connection; these are not obtained from the single complex order parameter and are deferred as an open problem in Section 14. The chain's claim at this step is therefore that stress-energy and the bosonic field sector are derivable; the extension to the full matter content is contingent on that open problem.

6.7. Step 8: from geometry and the vacuum to angular momentum

The final step in the chain is the one most often treated as an independent fact of nature: that massive systems orbit, that angular momentum is conserved, and that stable circular motion is ubiquitous from electrons to galaxies. In PST all of this follows from a single structural fact about the instantiated vacuum.

Once the order parameter is promoted to a complex field $\psi \in \mathbb{C}$, the vacuum manifold of the post-threshold system is $\mathcal{V} \cong S^1$, a one-dimensional circle of degenerate minima. The modal potential functional is invariant under global phase rotations $\psi \mapsto e^{i\alpha}\psi$: a continuous U(1) symmetry. By Noether's theorem, this symmetry implies a conserved current and a conserved charge. That charge is angular momentum in configuration space, $L = 2r_0^2\,\dot{\theta}$, and it projects via Π into the physical angular momentum of any instantiated system. The stable circular orbits of planets, the spin of black holes, the angular momentum of electrons in atomic orbitals, all reduce to the single fact that the instantiated vacuum is a circle and that motion along a circle is the only available ground state. Step 8 is therefore not an empirical input to PST; it is a derivation from the symmetry of the vacuum that itself follows from steps 1 through 4.

6.8. The chain as a closed logical structure

What is remarkable about the chain is not that it is long but that it is closed: no step requires a premise from outside the chain. Step 1 is the single primitive. Steps 2 through 8 follow from it by structural derivation, each one requiring only the concepts and structural commitments already established. No constants are introduced except as projective unit conversions. No laws are postulated except as structural theorems. No initial conditions are imposed except as consequences of the vacuum topology. Where the derivation relies on architectural choices, in particular the Landau-Ginzburg functional form and the gradient regularisation, these are the minimal choices consistent with the structural requirements; they are logically motivated rather than strictly entailed from (D, δ) alone, and this distinction is maintained explicitly throughout the paper.

This closure has a significant implication for the question of why the chain terminates at step 8 rather than continuing indefinitely. The chain does not terminate arbitrarily. Step 8 (angular momentum as Noether charge) is the last step for which the derivation is purely structural. Beyond step 8, physics enters the domain of the specific: particular particle masses, coupling constants, and quantum numbers that depend on the detailed content of the precausal configuration C rather than on the structural form of the chain. PST does not claim to derive the Standard Model of particle physics from first principles, any more than topology claims to derive the specific dimensions of any particular manifold. What PST claims is that the framework within which all such specific facts are intelligible, spacetime, causality, matter, and angular momentum, is not a brute given but a necessary structural consequence of a single irreducible primitive.

Other theories of physics typically either take this framework for granted or introduce it as an axiom without derivation. General Relativity assumes a Lorentzian manifold and derives its dynamics. Quantum mechanics assumes a Hilbert space, a Hamiltonian, and a measurement postulate, then derives its predictions. String theory assumes extra spatial dimensions and the string action. Loop quantum gravity assumes a discrete quantum geometry. Each of these frameworks contains ungrounded primitives: structures that are posited without explanation. The present chain shows that if one is willing to start from a single primitive that is itself logically irreducible, namely the capacity for distinctions to exist, all the background structure that modern physics takes as given follows as a necessary consequence. The chain is, in this sense, the logical spine of PST: the sequence of entailments that connects the most minimal possible ontology to the full landscape of physical structure.

From the perspective of the precausal substrate, there is one structure, the tension functional $T(C)$, projected through one operator Π into one manifold (M, g). Quantum behaviour and classical geometry are different aspects of the same projection, distinguished only by proximity to the modal threshold: near the threshold, the order parameter is small and the geometry is high-dimensional and rapidly varying, the quantum regime; far from the threshold, the order parameter is large and the geometry has settled into the four stable dimensions of the instantiated vacuum, the classical regime. The apparent gap between quantum mechanics and general relativity is, from this vantage point, not a gap between two fundamental theories but a gap between two limiting descriptions of a single underlying structure.

6.9. Layman's Summary

This section assembles all the preceding steps into a single logical chain of eight links: property differentiation, asymmetric tension, the modal threshold, modal sublimation, spacetime, causality, matter and energy, and finally angular momentum. Each link follows necessarily from the previous one, with no external input required. The chain is closed: remove any link and everything downstream collapses. What is remarkable is not its length but its completeness, the entire background structure of modern physics emerges from a starting point so minimal that no simpler one could be found.

7. Energy, Matter, and Fields

Einstein's field equations cast the relationship between matter and geometry in the form of a *source equation*: the Einstein tensor $G_{\mu\nu}$, which encodes the curvature of spacetime, is set equal to $8\pi G$ times the stress-energy tensor $T_{\mu\nu}$, which encodes the distribution of matter and energy [1, 10]. This formulation is extraordinarily successful as a description of the instantiated world. As an ontological account it is, however, incomplete in a specific and well-known way: it treats matter and geometry as two distinct kinds of thing, one of which acts as the source of the other. The left-hand side of the field equations, geometry, is described as marble, smooth and exact. The right-hand side, matter, is described as wood, coarse and phenomenological. Einstein himself remarked on this asymmetry as a deficiency, expressing the aspiration that a complete theory would derive the right-hand side from geometry rather than inserting it from outside. No such derivation was found within the framework of General Relativity, because within that framework, matter and geometry are genuinely independent inputs: the metric describes the arena; matter describes what moves through it.

PST dissolves this distinction at the foundational level. Both geometry and matter derive from the same precausal source, the asymmetric tension functional $T(C)$, preserved across the phase transition Φ and projected by the tension projection operator Π. The geometry of the instantiated manifold is the structural expression of $T(C)$ in the domain of spatial relations;

v9.9

the matter content is the energetic expression of the same $T(C)$ in the domain of physical densities. They are not two inputs; they are two outputs of a single projection. The marble and the wood are cut from the same precausal stone.

7.1. The stress-energy tensor as projected tension

In the isotropic, homogeneous approximation the precise statement is:

$$T_{\mu\nu}^{(\text{iso})}(x) = \Pi(T(C))(x) \cdot g_{\mu\nu}(x) \tag{24}$$

where $\Pi(T(C))(x)$ is the projection of the precausal tension functional onto the point $x \in M$, as defined in equation (19). This equation should be read carefully. It does not say that matter is made of geometry, nor that geometry is made of matter. It says that both are made of something prior to either: the asymmetric tension of the precausal configuration, which the functor Φ converts into a Lorentzian manifold and which the operator Π distributes as a scalar density across that manifold. The metric $g_{\mu\nu}$ provides the geometric structure; $\Pi(T(C))$ provides the scalar weight. The isotropic stress-energy tensor is the product of these two, a density on the manifold whose value at each point reflects how much tension content the precausal configuration contributed to that region of the instantiated geometry.

The components of $T_{\mu\nu}$ acquire their standard interpretations naturally in this picture. The energy density T_{00} is the local projection of the tension magnitude: regions where $T(C)$ was large before sublimation appear in the instantiated manifold as regions of high energy density. The momentum flux T_{0i} reflects the directional distribution of tension across configurations that realise nearby points. The pressure components T_{ij} encode the isotropic or anisotropic character of the tension distribution. None of these quantities is postulated; all are derived from the single precausal input $T(C)$ through the action of Π.

7.2. The field equations as a conservation statement

Einstein's field equations then read:

$$G_{\mu\nu} = 8\pi G \cdot \Pi(T(C)) \tag{25}$$

The standard reading of this equation is dynamical: matter curves spacetime, and the degree of curvature is set by the amount of matter. The PST reading is different and more fundamental: the field equations are a **conservation statement** expressing that the total asymmetric tension of the precausal configuration is faithfully preserved across the sublimation threshold, distributed simultaneously as curvature on the left-hand side and as stress-energy on the right. The equations do not govern a process; they describe what the functor Φ conserves. Curvature and stress-energy are not cause and effect; they are two names for the same preserved quantity, registered by two different instruments, the geometric instrument on the left and the energetic instrument on the right.

This reinterpretation connects with, and goes beyond, the thermodynamic approach of Jacobson [18] and the entropic gravity programme of Verlinde [17]. Both of these frameworks derive the field equations from non-geometric principles, treating them as equations of state rather than as fundamental laws. PST agrees with this spirit but locates the non-geometric substrate at a more foundational level: not in thermodynamic degrees of freedom within the instantiated manifold, but in the precausal tension structure that precedes the manifold itself. Where Jacobson derives Einstein's equations from entropy on holographic screens, PST derives them from the structure of Φ and Π. The two derivations converge on the same equations from different directions, suggesting that the equations themselves are more fundamental than either of their known derivations individually.

7.3. The F–S dichotomy

PST carries two distinct functionals that serve different mathematical roles, and a reader familiar with field theory should be warned at the outset that they are not interchangeable.

$\mathcal{F}[\psi, C]$ is the **modal free energy**: a Landau–Ginzburg potential defined over precausal configuration space. It is Euclidean in character: positive-definite in the equilibrium sector, with no temporal argument. Its stationary condition $\delta\mathcal{F}/\delta\psi = 0$ identifies equilibrium configurations; it governs the threshold, the vacuum manifold $\mathcal{V}$, and which mode is coherent. No time appears; no Noether theorem applies. $\mathcal{F}$ is not an action and must not be treated as one.

$S[g, \phi]$ is the **instantiated action**: a standard Lorentzian functional defined over spacetime histories in the instantiated domain $\mathcal{G}$. Its stationary points $\delta S/\delta g^{\mu\nu} = 0$, $\delta S/\delta\phi = 0$ are equations of motion. Noether's theorem, conservation laws, and all dynamical structure belong here. S is a *projected* object: it does not exist in the precausal substrate; it is what $\mathcal{F}$ becomes under the action of Π, with the Bernoulli measure $d\mu$ mapping to the covariant measure $\sqrt{-g}\, d^4x$ and the gradient term $c|\nabla_{\mathcal{C}}\psi|^2$ projecting to the kinetic sector of $\mathcal{L}$.

The relationship between $\mathcal{F}$ and S is the PST instance of the standard Euclidean-to-Lorentzian correspondence. A Landau–Ginzburg free energy is formally a Euclidean action: it is minimised rather than made stationary, integrated against a positive-definite measure, with no timelike direction. The standard route from Euclidean to Lorentzian is Wick rotation: $t \to -i\tau$ reverses the sign of the time-kinetic term and changes the signature from $(+, +, +, +)$ to $(-, +, +, +)$. In PST, sublimation *is* the Wick rotation. The moment the $(-, +, +, +)$ Lorentzian signature emerges from the functor Φ is precisely the moment $\mathcal{F}$ ceases to be the relevant functional and S takes over. This is not an embarrassing gap in the formalism. It is a structural prediction: the Lagrangian cannot exist before the threshold because the time one would integrate it over does not yet exist. That S has no precausal counterpart is a theorem of the framework, not a shortcoming.

F–S Dichotomy. $\mathcal{F}$ governs the precausal domain: equilibrium, threshold, and vacuum structure, with no dynamics and no Noether charges. S governs the instantiated domain: equations of motion, conservation laws, and all of quantum field theory. The passage from $\mathcal{F}$ to S

is sublimation; it is not reversible; and nothing on the S-side can be traced back to the precausal level without crossing the threshold. Every subsequent use of Noether's theorem, every conserved charge, and every field-theoretic result in this paper belongs exclusively to S.

7.4. The projected action and the emergence of the Einstein–Hilbert term

The field equations (25) were presented above as a conservation statement, but their specific tensor structure: why the left-hand side is $G_{\mu\nu}$ and not $R_{\mu\nu}$, or $R\,g_{\mu\nu}$, or some other curvature combination: requires a derivation that has so far been deferred. That derivation begins with the projected action.

The potential functional $\mathcal{F}(\psi, C)$ of equation (33) contains three terms. The first two, $a|\psi|^2$ and $b|\psi|^4$, determine the equilibrium structure of the order parameter. The third, $c\,|\nabla_{\mathcal{C}}\psi|^2$, is the gradient term: it penalises variation of the order parameter across configuration space and controls the energetic cost of excitations above the vacuum. Under the projection Φ, this term acquires a precise geometric interpretation.

When the order parameter has settled to its vacuum value $\psi = r_0\,e^{i\theta}$ on $\mathcal{V}$, the configuration-space gradient $\nabla_{\mathcal{C}}\psi$ decomposes into two pieces: the Goldstone mode $r_0\,\nabla_{\mathcal{C}}\theta$ (responsible for matter and angular momentum, discussed below) and a *metric fluctuation* component, which encodes how the instantiated geometry responds to variations of the realisation map ρ. Integrating the metric fluctuation component over $\mathcal{C}$ with the measure μ and projecting onto (M, g) via Π, the leading term in a derivative expansion of the projected gradient is:

$$\Pi\big(c\,|\nabla_{\mathcal{C}}\psi|^2\big)\,\big|_{\psi=r_0} \;\longrightarrow\; \frac{r_0^2\,c}{16\pi G}\int_M R\,\sqrt{-g}\;\mathrm{d}^4x \;+\; \mathcal{O}\big(R_{\mu\nu}R^{\mu\nu},\,R^2\big) \tag{26}$$

where R is the Ricci scalar of (M, g). The coefficient $1/(16\pi G)$ sets the gravitational coupling scale and identifies G as derived from r_0, c, and the geometry of the projection: consistent with equation (30). The higher-order terms in $R_{\mu\nu}R^{\mu\nu}$ and R^2 are suppressed by $(k\,d_0)^2$ and are negligible at all scales $\gg d_0 \approx 7$ nm. The effective projected action at low energies is therefore:

$$S[g] \;=\; \frac{1}{16\pi G}\int_M R\,\sqrt{-g}\;\mathrm{d}^4x \;+\; S_{\text{matter}}[g,\,\text{fields}] \tag{27}$$

This is the **Einstein–Hilbert action**, emerging from the configuration-space gradient term in $\mathcal{F}$ without being postulated. Varying equation (27) with respect to $g^{\mu\nu}$ yields, by a standard calculation:

$$\frac{\delta S}{\delta g^{\mu\nu}} \;=\; 0 \qquad\Longrightarrow\qquad G_{\mu\nu} \;=\; 8\pi G\,T_{\mu\nu} \tag{28}$$

where $T_{\mu\nu} = -2/\sqrt{-g}\cdot\delta S_{\text{matter}}/\delta g^{\mu\nu}$ is the stress-energy tensor of the projected matter fields. The Einstein tensor $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R\,g_{\mu\nu}$ appears specifically because it is what the variation of $\int R\sqrt{-g}$ with respect to $g^{\mu\nu}$ produces, up to boundary terms. Equation (25) is thereby derived, not asserted.

7.5. Uniqueness of $G_{\mu\nu}$: diffeomorphism invariance and Lovelock's theorem

A further question remains: why is the Einstein–Hilbert action the correct projected action, rather than one involving $R_{\mu\nu}R^{\mu\nu}$, or R^2, or other higher-curvature terms? The answer is a two-step argument using the diffeomorphism invariance that PST derives in Section 4.

Step 1: Noether's second theorem. The projected action $S[g]$ inherits diffeomorphism invariance from the derivation of Section 4: any two choices of realisation map ρ produce diffeomorphic instantiated geometries, so $S[g]$ is invariant under the infinite-dimensional group $\mathrm{Diff}(M)$. By Noether's second theorem [20], invariance of an action under a local (gauge) symmetry group implies that the Euler–Lagrange expressions satisfy differential identities: the Noether identities: holding off-shell, as an algebraic consequence of the symmetry, not only on solutions. For diffeomorphism invariance the Noether identity reads:

$$\nabla^\mu \left(\frac{\delta S}{\delta g^{\mu\nu}} \right) = 0 \tag{29}$$

This forces the left-hand side of the field equations to be divergence-free: $\nabla^\mu G_{\mu\nu} = 0$. Applied simultaneously to the matter sector, equation (29) gives $\nabla^\mu T_{\mu\nu} = 0$: stress-energy conservation follows from PST-derived diffeomorphism invariance, not from a separate postulate.

Step 2: Lovelock's theorem. The tensor on the left-hand side of the field equations must be: (i) symmetric, since it arises by varying a diffeomorphism-invariant scalar action with respect to the symmetric tensor $g^{\mu\nu}$; (ii) divergence-free, by the Noether identity (29); and (iii) built from $g_{\mu\nu}$ and its first and second derivatives alone, since the projected gradient term generates at most second-order curvature (the $\mathcal{O}(R_{\mu\nu}R^{\mu\nu})$ terms of equation (26) are negligible in the low-energy effective theory). Lovelock's theorem [21] establishes that in four spacetime dimensions, the *unique* symmetric, divergence-free, second-rank tensor built from $g_{\mu\nu}$ and at most its second derivatives is $\alpha\, G_{\mu\nu} + \Lambda\, g_{\mu\nu}$ for constants α and Λ. Condition (iii) is what places the projected action inside the scope of this theorem: PST derives the second-order structure from the suppression of higher-curvature terms, so invoking Lovelock is not an additional assumption but a consequence of the projected-gradient derivation. With $\alpha = 1/(16\pi G)$ from equation (27) and Λ identified with the cosmological constant, the Einstein tensor is the unique second-order result. More precisely: *within the low-energy effective theory defined by the projected second-order gradient*, $G_{\mu\nu} + \Lambda g_{\mu\nu}$ is the only admissible tensor structure. This is the appropriate effective-theory sense of uniqueness; it does not preclude higher-curvature corrections at scales $k \sim 1/d_0$, where the suppressed $\mathcal{O}(R_{\mu\nu}R^{\mu\nu})$ terms of equation (26) become relevant. This answers why gravity takes the specific form it does at observable scales: PST derives the diffeomorphism invariance that forces the divergence-free condition; Lovelock's theorem then closes the argument to uniqueness within the second-order effective action.

7.6. The equivalence principle as a theorem of single-source projection

General Relativity postulates the weak equivalence principle: that all matter falls identically in a gravitational field, regardless of composition: as an empirical fact encoded in the theory but not explained by it. In PST, the equivalence principle is not postulated. It is a structural consequence of the fact that all matter is a projection of a single source through a single operator.

Every matter species, whatever its internal modal structure, is a projection of the precausal tension functional $T(C)$ through the single operator Π. The coupling between any matter field and the instantiated geometry is mediated exclusively by Π, which carries no species label. The gravitational coupling constant $8\pi G$ is therefore the same for every projected matter field, not by assumption, but because there is only one projection operator and one source. There is nothing for different matter species to couple differently *to*.

Formally: the geodesic equation for any test body follows from stress-energy conservation $\nabla^\mu T_{\mu\nu} = 0$, derived from the Noether identity (29). That derivation makes no reference to the body's composition or rest mass; it uses only the fact that the body's stress-energy is a projection of $T(C)$ through Π. The free-fall trajectory is therefore universal: a geodesic of (M, g), identical for all bodies of negligible self-gravity. The weak equivalence principle is a theorem of single-source projection.

7.7. The gravitational constant as a unit conversion

The factor $8\pi G$ in equation (25) deserves separate attention. In standard physics, Newton's gravitational constant G is a dimensionful fundamental constant whose numerical value is fixed by measurement and whose theoretical derivation is unknown. In PST, G is not a property of the precausal substrate. It is a property of the projection operator Π: the factor that converts between dimensionless modal tension units and physical stress-energy units. It enters the field equations as a unit conversion, not as a coupling strength.

The significance of this reframing is substantial. If G is a fundamental constant, then a quantum theory of gravity must quantize it, and the question of why G has the value it does becomes a question about the fine structure of nature with no obvious answer. If G is a unit conversion introduced by Π, then its value is set by the structure of the projection map, a structure that is in principle derivable once the full form of Π is specified. The dimensionless quantities that ultimately determine the ratio of gravitational to electromagnetic effects are properties of the configuration C, not of the substrate structure. This is consistent with the Dirac large numbers hypothesis [45] and with the programme of deriving coupling constants from deeper structural relationships, though PST locates the derivation in the precausal domain rather than in a numerological coincidence.

7.8. Derivation of G from the projection scale

The preceding paragraph identified G as a unit conversion factor. This identification is the beginning of a derivation, not the end of it. To complete it, one must determine what structural feature of Π fixes the numerical value of G.

In natural units where $\hbar = c = 1$, Newton's constant carries dimensions of $[\text{length}]^2$: one has

$G = \ell_P^2$, which defines the Planck length $\ell_P \approx 1.616 \times 10^{-35}$ m as the length unit in which $G = 1$. The projection operator Π of equation (19) has kernel $K(x, C') = \delta(x - \rho(C'))$, which localises the tension of each precausal configuration C' at the spacetime point $x = \rho(C')$. The realization map ρ carries a characteristic length: the smallest separation $d_0 > 0$ at which distinct properties $a, b \in D$ are assigned distinguishable spacetime points. Below d_0, the discreteness of the distinction space D is manifest and the continuum description breaks down; above d_0, the projection produces a smooth geometry well-approximated by a Lorentzian manifold. The scale d_0 is therefore the intrinsic grain size of the precausal structure, the finest resolution at which the instantiated geometry can represent distinctions from D.

The unit conversion performed by Π involves d_0 in a specific and computable way. The delta-function kernel $\delta(x - \rho(C'))$ carries units of $[\text{length}]^{-d}$ where d is the number of spatial dimensions of M. When the integral over $\mathcal{C}$ in equation (19) is evaluated, the measure $\mu(C')$, dimensionless as derived in Section 3 via the canonical Bernoulli measure of equation (10), contributes no length scale; all dimensional content enters through the kernel alone. In four spacetime dimensions, matching the dimensions of $\Pi(T(C))(x)$ to those required by the field equation (25) yields a dimensionless coefficient α_G fixed by the detailed form of K:

$$G \;=\; \alpha_G \cdot d_0^2 \qquad (\hbar = c = 1) \tag{30}$$

Since $G = \ell_P^2$ in these units, equation (30) states that $d_0 = \alpha_G^{-1/2}\, \ell_P$. Crucially, the equation does not constrain α_G: it *defines* $\alpha_G \equiv G/d_0^2$ and defers the question of d_0's actual magnitude to a separate derivation. Section 10 carries out that derivation via PST's dimensional reduction mechanism, showing that d_0 is not the Planck length but the coherence length of the modal condensate, fixed at approximately 7 nm by the electroweak modal scale, twenty-six orders of magnitude above ℓ_P and within reach of laboratory Casimir measurements. The Planck length marks the scale at which classical geometric descriptions break down from dimensional arguments involving $\hbar$, c, and G; d_0 is a distinct, physically derivable scale at which the continuum approximation of Π ceases to hold.

The consequence is immediate. The question of determining d_0 from the structure of (D, δ, T) and the question of deriving G from Π are not independent: equation (30) shows they are the same question. A determination of d_0 automatically delivers G, with α_G computable from the geometry of K. What the equation achieves is a *reduction*: two seemingly independent open problems collapse into one. Section 10 shows that the dimensional reduction inherent in modal sublimation fixes d_0 conditionally on identifying the fundamental modal scale ℓ_* with the electroweak condensate, yielding $d_0 \approx 7$ nm and $\alpha_G \approx 5 \times 10^{-54}$. The unconditional derivation of ℓ_* from (D, δ, T) alone remains an open problem (Section 14, Problem 1).

7.9. The cosmological constant reinterpreted

The cosmological constant Λ enters the field equations as a constant background energy density of the vacuum, contributing a term $\Lambda g_{\mu\nu}$ to the left-hand side [19]. Standard quantum field theory predicts a vacuum energy density from zero-point fluctuations of all quantum fields, but the predicted value exceeds the observed Λ by approximately 120 orders of magnitude, a discrepancy that has been called the most embarrassing problem in theoretical physics.

PST dissolves this problem by identifying the physical vacuum not with $\psi = 0$ but with the ring of degenerate minima $\mathcal{V} \cong S^1$ at radius r_0. The zero-point fluctuation calculation in QFT implicitly sums all field modes at the symmetric vacuum $\psi = 0$, which is the unstable maximum of the sombrero potential in the instantiated regime. This is not the physical vacuum; it is the energy of the wrong ground state. The correct vacuum energy is the value of the potential functional $\mathcal{F}$ evaluated on $\mathcal{V}$:

$$\mathcal{F}\big|_{\mathcal{V}} = \int_{\mathcal{C}} \left[-\frac{(T(C) - \tau)^2}{4b} \right] d\mu(C) \tag{31}$$

This is negative (the sombrero valley lies below the origin) and finite, set by the excess tension $\varepsilon(C) = T(C) - \tau$ integrated over all instantiated configurations. The cosmological constant is therefore not a free parameter but a derived quantity, and the standard QFT calculation is an artefact of evaluating the vacuum energy at the wrong point in field space.

7.10. The sign of the cosmological constant: a quantitative tension

The valley energy above is *negative*: the sombrero ring lies below the symmetric maximum. A negative vacuum energy density sources Anti-de Sitter space: constant negative curvature, decelerating, recollapsing. The observed cosmological constant, by contrast, is positive: the universe is in accelerating, de Sitter-like expansion. This is a genuine tension inside the PST gravity treatment and is acknowledged here without evasion.

The PST framework supplies a candidate mechanism for a sign reversal. The vacuum is not static: configurations on $\mathcal{V}$ carry the Noether charge $L = 2r_0^2\dot{\theta}$ of the spontaneously broken U(1) symmetry (derived in the angular-momentum paragraph below). This orbital motion contributes a positive kinetic energy to the effective vacuum energy density:

$$\rho_{\text{vac}} = \underbrace{\frac{L^2}{8r_0^2}}_{\text{orbital kinetic (positive)}} + \underbrace{\mathcal{F}\big|_{\mathcal{V}}}_{\text{valley depth (negative)}} \tag{32}$$

De Sitter expansion requires $\rho_{\text{vac}} > 0$, i.e. the kinetic term must exceed the valley depth: $L^2/(8r_0^2) > |\mathcal{F}|_{\mathcal{V}}|$. A second candidate mechanism operates at the boundary of instantiation: configurations approaching the threshold from below exert an effective outward tension gradient on the instantiated domain, sourcing a positive pressure contribution from ongoing sublimation. The sign of this gradient is fixed by the direction of threshold crossing (tension increasing toward τ) and is always positive.

Neither mechanism is quantitatively closed by the current formalism. Deriving L independently of G and the observed $\Lambda \approx 1.1 \times 10^{-52}\ \text{m}^{-2}$, and showing that equation (32) reproduces the correct magnitude, remains an open problem (Section 14, Problem $2a$). The tension is real;

the mechanism for its resolution is native to PST; the quantitative work is deferred.

7.11. The Degenerate Vacuum Manifold and the Origin of Angular Momentum

The treatment of the order parameter ψ in Section 3 took ψ to be a real scalar. This is sufficient to establish the bifurcation and derive the threshold condition. It is not, however, sufficient to capture the full modal structure of the instantiated vacuum. A precausal configuration possesses two independent components of modal imbalance: the magnitude of the tension $T(C)$ and its orientation within configuration space, reflecting the fact that asymmetric tension is directional. These two components are correctly represented by treating the order parameter as a complex scalar $\psi \in \mathbb{C}$, written in polar form as $\psi = r\, e^{i\theta}$, with $r \geq 0$ and $\theta \in [0, 2\pi)$.

Under this extension, the potential functional becomes:

$$\mathcal{F}(\psi, C) = \int_{C} \Big[-\varepsilon\, |\psi|^2 + b\, |\psi|^4 + c\, |\nabla_{\mathcal{C}} \psi|^2 \Big]\, d\mu(C) \tag{33}$$

where $|\psi|^2 = \psi_1^2 + \psi_2^2$ and $\psi = \psi_1 + i\psi_2$. This functional is invariant under global phase rotations $\psi \mapsto e^{i\alpha}\psi$ for all $\alpha \in \mathbb{R}$: a continuous U(1) symmetry. For $\varepsilon(C) > 0$, the unique minimum is $\psi = 0$, the symmetry is unbroken, and the potential is a paraboloid, a single symmetric bowl with no preferred phase. For $\varepsilon(C) < 0$, the symmetric minimum becomes an unstable saddle point, and the functional acquires a continuous ring of degenerate minima:

$$\mathcal{V} = \big\{\, \psi \in \mathbb{C} \;:\; |\psi| = r_0 \,\big\}, \qquad r_0 = \sqrt{\frac{T(C) - \tau}{2b}} = \sqrt{\frac{\varepsilon}{2b}} \tag{34}$$

The vacuum manifold $\mathcal{V}$ is topologically S^1. In the two-component representation (ψ_1, ψ_2), the potential surface takes the form of a *sombrero potential*: a local maximum at the origin surrounded by a continuous circular valley of degenerate minima at radius r_0. It is this structure, not the paraboloid, that constitutes the correct vacuum of any instantiated configuration.

The spontaneous breaking of the U(1) symmetry, when the system selects a specific phase θ_0 and settles into one point on $\mathcal{V}$, is accompanied by a massless excitation in the θ direction, the Goldstone mode of the broken symmetry [44]. In the precausal context, this massless mode corresponds physically to motion along $\mathcal{V}$ at zero energetic cost: the phase of the order parameter can vary without climbing the potential, because all points on the ring are degenerate. After projection through Π, this massless precausal excitation manifests as a massless propagating mode in the instantiated geometry.

Before making the identification with the photon, a structural caveat is required. The U(1) symmetry broken here is a *global* symmetry of a configuration-space scalar. The U(1) gauge symmetry of electromagnetism is a *local* symmetry of a spacetime vector field. The mathematical move from global to local symmetry is *gauging*: introducing a connection on a principal U(1) bundle over M, and promoting ordinary derivatives to gauge-covariant ones. This step produces the photon field as the connection, and is responsible for the minimal coupling of charged matter to the electromagnetic field. The gauging step is deferred to future work and is listed as an open problem in Section 14. Subject to this caveat, the identification of the Goldstone mode with the photon is structurally motivated: both are massless modes of a broken U(1) symmetry, the electromagnetic U(1) is the gauged version of exactly this global symmetry, and the observed masslessness of the photon is consistent with this origin.

7.12. The incomplete vacuum of standard physics

The vacuum formulation underlying General Relativity and its semiclassical extensions assumes a unique, symmetric ground state. The cosmological constant Λ enters Einstein's field equations [19] as a constant energy density associated with the symmetric minimum $\psi = 0$. This is the $\varepsilon > 0$ regime of the PST potential, the pre-threshold paraboloid, and it correctly describes configurations that have not yet crossed the modal threshold. Once sublimation has occurred and $\varepsilon < 0$, the symmetric minimum is no longer a stable state; it is the local maximum at the centre of the sombrero. The energy assigned to $\psi = 0$ in the standard treatment is the energy of an unstable critical point, not a ground state.

The assumption of a unique symmetric vacuum is exact within its domain: for pre-threshold configurations, the paraboloid is the correct potential. The error is one of domain. The cosmological vacuum is an instantiated vacuum, and instantiated configurations inhabit the $\varepsilon < 0$ regime. The physical ground state is not $\psi = 0$ but any point on $\mathcal{V} \cong S^1$, selected by spontaneous breaking of the U(1) symmetry. What no theory built on the assumption of a unique symmetric vacuum can account for, without additional postulate, is why every gravitationally bound system in the observable universe, from the orbits of planets to the angular momentum of rotating black holes, exhibits stable circular structure. PST requires no additional postulate.

7.13. Angular momentum as Noether charge

The global U(1) symmetry of $\mathcal{F}$ under $\psi \mapsto e^{i\alpha}\psi$ is inherited by the instantiated action S through the projection Π: the Goldstone mode $\theta(\mathbf{x})$ enters S as a massless spacetime scalar with a standard kinetic term, and the phase rotation is a global symmetry of that term. Noether's theorem [20] applies to S, not to $\mathcal{F}$ directly; $\mathcal{F}$ is a free energy whose minima $\mathcal{V}$ determine which symmetry is spontaneously broken, while S is the Lorentzian action whose dynamics generate the conserved charge. This distinction matters: Noether charges require a time for the charge to be conserved *along*; time is coinstantiated, so of course conservation laws are instantiated-domain structures.

With $\psi = r\, e^{i\theta}$ and the Goldstone kinetic term $r_0^2\, \partial_\mu\theta\, \partial^\mu\theta$ in $\mathcal{L}$, the Noether current associated with the phase rotation is:

$$j^\mu = 2r^2\, \partial^\mu\theta \tag{35}$$

where ∂^μ are spacetime derivatives in $\mathcal{G}$. The conserved Noether charge is:

$$L = \int j^0\, d^3x = 2r_0^2\, \dot{\theta} \tag{36}$$

where $\dot{\theta}$ is the time derivative of the phase in the instantiated domain. This charge L is angular momentum. It is not postulated; it is the Noether charge of the projected instantiated action corresponding to the U(1) symmetry that $\mathcal{F}$ breaks spontaneously at threshold.

Before making the identification with physical angular momentum, a distinction must be drawn between two related but separate claims. The first is that the spontaneous breaking of U(1) symmetry generates a conserved Noether charge L in configuration space. This is established above and is rigorous. The second is that this charge projects via Π into the macroscopic orbital angular momentum of gravitationally bound physical systems. These claims are related but not identical. In standard field theory, the Goldstone mode of a broken U(1) is a massless field excitation propagating through the geometry; it is not automatically identified with the macroscopic orbital motion of massive bodies. The identification of L with physical orbital angular momentum therefore requires the explicit action of Π: the projection of the configuration-space Noether charge onto the instantiated geometry must be shown to recover the orbital angular momentum of bodies moving on geodesics. The structural argument for this identification is clear, as Π is by construction the structure-preserving map between the precausal and instantiated domains, and any conserved charge in configuration space must project to a conserved quantity in the instantiated geometry. However, the explicit computation of this projection, showing in detail how $L = 2r_0^2\dot{\theta}$ maps to the orbital angular momentum of a planetary body or an electron in an atomic orbital, is deferred to future work and is listed as an open problem in Section 14.

With that qualification in place: in the instantiated domain, L projects via Π into physical angular momentum. Motion along the vacuum manifold, the order parameter tracing the ring of degenerate minima, is, after projection, the circular orbital motion of any instantiated system. Orbital mechanics is not an initial condition imposed from outside; it is the Noether charge of the spontaneously broken U(1) symmetry of the vacuum, conserved by construction. The stable circular orbits of gravitationally bound systems, which General Relativity describes correctly but does not derive, are, in PST, the structural projection of U(1) symmetry conservation across the sublimation threshold, subject to the derivation noted above.

The quantisation of L in the near-threshold regime deserves comment. Near the modal threshold, where $\varepsilon(C) \approx 0$, the order parameter ψ is not settled at a definite point on $\mathcal{V}$ but fluctuates across the ring. The charge L is then not a classical continuous quantity but takes discrete values, the Noether charge computed over a fluctuating path on S^1. The quantisation condition $L = n\hbar$ for integer n is the precausal statement that the phase of the order parameter can only wind an integer number of times around the ring without discontinuity. Spin is the near-threshold manifestation of the same U(1) charge that produces orbital angular momentum in the far-threshold classical regime.

7.14. The membrane analogy and its implicit error

The standard pedagogical representation of General Relativity, the rubber sheet or elastic membrane stretched by the mass of an embedded object, depicts spacetime curvature as a funnel-shaped depression: a single minimum toward which neighbouring objects roll [10, 22]. This image is vivid and heuristically powerful, but it encodes precisely the incomplete vacuum identified above. A funnel has one minimum. In the language of the PST potential, it is a paraboloid, the $\varepsilon > 0$ pre-threshold potential. An object placed anywhere on this surface rolls toward the centre and stops there; the only way to produce a stable orbit in the membrane model is to give the object a precisely tuned initial tangential velocity, imposed by hand as an external condition with no explanation from the geometry itself. The membrane is silent on where this velocity comes from. It can describe an orbit once one is assumed, but it cannot derive the existence of orbital motion from the structure of the vacuum.

The deficiency is not merely pedagogical. It is present in the geodesic framework itself. In General Relativity, the geodesic equation correctly determines every possible path through a given metric, but selecting which geodesic a body actually follows remains an initial condition. The geometry provides a catalogue of paths; it does not explain why any particular one is occupied. This is a foundational gap: the theory is kinematically complete and dynamically incomplete with respect to the question of why bodies move at all rather than remaining stationary.

The sombrero potential removes this deficiency without remainder. In the $\varepsilon < 0$ regime, the centre of the funnel is not a minimum but an unstable maximum: any configuration placed there is immediately expelled outward. The minimum is the ring $\mathcal{V}$, and motion along this ring, the order parameter tracing S^1, is the unique stable ground state of the vacuum. An object coupled to this vacuum does not orbit *despite* the geometry; it orbits *because* the geometry has no other resting place at the level of the global vacuum. The circular structure of gravitationally bound systems, planetary orbits, galactic rotation, the spin of black holes, is, in this picture, not a mechanical accident but a direct reflection of the topology of the vacuum manifold.

This argument applies to systems that are coupled to the global U(1) Noether charge of the vacuum. It should not be read as a claim that every stable configuration in the universe must exhibit circular motion. Bound states such as molecular bonds and crystal lattices represent local potential minima that break translational symmetry without requiring the global U(1) charge; they are stable against the local chemical potential rather than the global vacuum topology. The sombrero argument addresses the gravitational and electromagnetic regimes, where the U(1) conservation law is operative, and explains why those regimes exhibit universal circular structure. It does not apply uniformly to all scales and all interactions, and no such universal claim is made.

7.15. The projected field theory: emergence of quantum fields

The Einstein–Hilbert term was extracted from the gradient term $c|\nabla_{\mathcal{C}}\psi|^2$ by evaluating it on the settled vacuum $\psi = r_0\, e^{i\theta_0}$: the metric-fluctuation branch of the gradient decomposition carried the geometric content. That decomposition has a second branch. The same gradient term also governs fluctuations of the order parameter *about* the vacuum, and it is this branch, not the vacuum value but the spectrum of excitations around it, that projects into quantum field theory. Geometry is what the gradient term contributes at the vacuum; fields are what it contributes away from it.

Writing the order parameter near a point on $\mathcal{V}$ as

$$\psi(\mathbf{x}) = \big(r_0 + h(\mathbf{x})\big)\, e^{i(\theta_0+\theta(\mathbf{x}))}, \tag{37}$$

and retaining terms through quadratic order in the fluctuations h and θ, the gradient term gives $(\nabla h)^2 + (r_0 + h)^2(\nabla\theta)^2$, and the potential $a|\psi|^2 + b|\psi|^4$ contributes a curvature in the radial direction and nothing in the phase direction (the flat direction along $\mathcal{V}$). The radial curvature at r_0, with $r_0^2 = \varepsilon/(2b)$ from equation (34), is

$$m_h^2 = \mathcal{F}''(r_0) = 4\varepsilon, \tag{38}$$

so the radial mode is massive with a mass set by the excess tension $\varepsilon = T(C) - \tau$, while the phase mode is exactly massless. This is the precise structural realisation of the qualitative claim that mass measures distance from the threshold: deeper instantiation (ε larger) produces a heavier radial excitation, while the Goldstone phase mode sits at zero mass because motion along $\mathcal{V}$ costs nothing.

The fluctuation functional is at this stage Euclidean, inherited from $\mathcal{F}$: positive-definite, no timelike direction, integrated against the Bernoulli measure. It is not yet a field theory in the dynamical sense. The functor Φ supplies what is missing. As established by the F–S Dichotomy above, sublimation is the PST instance of Wick rotation: the moment the $(-,+,+,+)$ signature emerges from Φ is the moment the Euclidean quadratic functional becomes a Lorentzian quadratic action,

$$S_2[h,\theta] = \int_M \left[\tfrac{1}{2}(\partial h)^2 - \tfrac{1}{2}m_h^2\, h^2 + \tfrac{1}{2}r_0^2\,(\partial\theta)^2\right]\sqrt{-g}\; d^4x + \dots, \tag{39}$$

where the ellipsis denotes interaction terms from $b|\psi|^4$. The radial mode obeys a massive Klein–Gordon equation; the phase mode obeys a massless one. Their propagators are the inverses of the respective quadratic operators, with the $i\varepsilon$ prescription not imposed by hand but fixed by the Lorentzian signature that Φ produces. Poincaré invariance of S_2 in the Minkowski limit is not a separate postulate: it is the flat-space restriction of the diffeomorphism invariance derived in Section 4. The projected field theory is automatically relativistic for the same reason the projected gravity is generally covariant.

Two distinct discrete structures arise once the field theory exists, and they must not be conflated. The first is topological. The vacuum manifold $\mathcal{V} \cong S^1$ has fundamental group $\pi_1(S^1) = \mathbb{Z}$, so the phase θ can wind around $\mathcal{V}$ only an integer number of times: a configuration cannot return to itself across a fractional winding without discontinuity. This integer is the U(1) charge, and it is the origin of the discrete charge spectrum $L = n\hbar$ derived in the vacuum-manifold discussion above, angular momentum in the far-threshold regime, spin near it. This quantisation is genuinely structural: it follows from the topology of $\mathcal{V}$ alone, with no quantisation postulate imposed.

The second discrete structure is the Fock space, and it is a separate matter. The occupation number of field quanta is the eigenvalue of $N = a^\dagger a$, an integer unrelated to the topological winding number: the former counts excitations of the projected field, the latter labels superselection sectors of the phase. The Fock space arises in PST exactly as in ordinary field theory, by canonical quantisation of the projected Lorentzian action S_2 once it exists, that is, once sublimation has supplied the timelike direction along which equal-time commutation relations can be posed. PST therefore derives the arena in which canonical quantisation takes place and the charge spectrum the resulting states must carry, but it does not yet derive the canonical commutation relations themselves from (D, δ, T); that derivation is part of the same programme as the Born rule and is listed together with it in Section 14. What the present construction establishes is narrower and should be stated as such: the projected field theory is a standard Lorentzian field theory carrying a topologically quantised U(1) charge, and its excitation spectrum is discrete for the ordinary reason, not because winding number and occupation number coincide.

The projection operator Π carries the intrinsic scale d_0 derived in Section 10. Below d_0 the continuum description of Π fails; above it the smooth field theory holds. The scale $1/d_0$ is a physical ultraviolet cutoff, not a calculational regulator. The divergences of standard QFT, recovered as $d_0 \to 0$, are artefacts of taking the continuum limit of a projection that was never continuous at arbitrarily short distances; renormalisability is the statement that low-energy observables are insensitive to d_0, which holds for separations $\gg 7$ nm precisely as the Casimir analysis of Section 10 requires.

The following items are not yet supplied by this construction and are stated explicitly:

The Born rule. The identification of $|\psi|^2$ near threshold with measurement probability is structurally motivated but not derived from $\mathcal{F}$. Derivation is listed as an open problem in Section 14.

Spin and statistics. The excitations constructed here are bosonic, carrying integer winding charge. Half-integer spin, the fermionic sector, and the spin–statistics connection require structure beyond the single complex order parameter and are not obtained by the present construction. This is a genuine gap and is listed as an open problem in Section 14.

The full interacting theory. The free-field Fock space and the ϕ^4-type vertices are exhibited; a rigorous treatment of the interacting vacuum and its renormalisation is deferred.

The relationship between PST and quantum field theory is now exactly parallel to its relationship with General Relativity, and the parallel is the point. QFT is neither a competitor to PST nor identical to it. It is the Lorentzian projection, through Π, of the fluctuation spectrum of $\mathcal{F}$ about $\mathcal{V}$, just as General Relativity is the projection of the vacuum value of the same functional. PST is not itself a quantum field theory and could not be: $\mathcal{F}$ is Euclidean, atemporal, and precausal, and a quantum field theory requires the Lorentzian action S that exists only after the threshold is crossed. To demand that PST be a QFT at the foundational level would be to demand that it presuppose the spacetime, causal ordering, and time-ordered structure it sets out to derive. That QFT has no precausal counterpart is, as with the absence of a precausal action S, a theorem of the framework rather than a shortcoming.

7.16. Layman's Summary

This section derives Einstein's field equations as consequences rather than postulates. The stress-energy tensor, which encodes the distribution of matter and energy, turns out to be the same precausal tension that produced the spacetime geometry, now expressed in geometric units. Gravity and matter are not two separate things connected by the field equations; they are two faces of one underlying structure. The section also shows that the vacuum (the ground state of the universe after sublimation) has the shape of a ring rather than a single point in an internal sense, and that motion along this ring is, by a symmetry theorem, the structural origin of all orbital motion in the universe, from atomic electrons to planetary orbits to galaxies.

8. Gauge Structure, Quantization, and the Standard Model

The projected field theory of Section 7 establishes a Lorentzian scalar field theory carrying a topologically quantised U(1) charge, with a massive radial mode (the Higgs) and a massless Goldstone mode. That treatment was explicit about two structural gaps: the step from global to local gauge invariance was deferred, and the non-Abelian factors $\mathrm{SU}(2)_L \times \mathrm{SU}(3)_c$ of the Standard Model were expected but not derived. The present section closes both gaps and establishes a third result: gauge anomaly cancellation is not a numerical coincidence of the Standard Model but a structural theorem whose proof rests on PST's derivation of $N = 3$ compact substrate dimensions in Section 11. Together, these results constitute the gauge-theoretic skeleton of a complete PST derivation of the Standard Model.

8.1. Path-integral quantization of the projected theory

The Euclidean modal functional $\mathcal{F}[\psi, C]$, equipped with the Bernoulli product measure μ derived in Section 3, defines a substrate partition function:

$$Z_{\mathrm{sub}} \;=\; \int_{\mathcal{C}} \exp\!\big(-\mathcal{F}[\psi(C), C]\big)\, d\mu(C). \tag{40}$$

The functor $\Phi : \mathbf{S} \to \mathbf{G}$ maps each configuration to an instantiated spacetime geometry; the F–S dichotomy of Section 7 establishes that this is the precausal counterpart of Wick rotation. The instantiated partition function is therefore:

$$Z_{\mathrm{QFT}} \;=\; \int \mathcal{D}\psi\, \mathcal{D}g_{\mu\nu}\, \exp\!\big(iS[\psi, g]\big), \tag{41}$$

with $S = S_{\mathrm{EH}} + S_{\mathrm{matter}}$ the Einstein–Hilbert plus matter action of Section 7. This is the path integral of quantum gravity coupled to the modal field. The substrate grain d_0 derived in

Section 11 provides a physical ultraviolet cutoff $\Lambda_{\mathrm{UV}} = 1/d_0$: not a calculational regularisation but a genuine breakdown scale of the continuum description. Renormalisability is the statement that low-energy observables are insensitive to d_0, which holds for separations $\gg 7$ nm exactly as the Casimir analysis requires.

Expanding around the vacuum $\psi = r_0 e^{i\theta_0}$, the quadratic fluctuation action S_2 of equation (39) generates Feynman rules for a massive scalar (mass $m_h^2 = 4\varepsilon$, equation (38)) and a massless pseudoscalar (the Goldstone mode), together with graviton propagators from the metric fluctuations $h_{\mu\nu}$. The non-Abelian gauge bosons enter through the extended order-parameter structure of the paragraphs below.

8.2. From global to local gauge symmetry: the photon as a connection

The global U(1) symmetry $\psi \mapsto e^{i\alpha}\psi$ is a symmetry of the action S_2 with α a single number. But the projection Π maps the continuum of substrate configurations to the full manifold (M, g): different loci of M are images of different configurations C, each of which independently selects a vacuum phase $\theta(x) \in \mathcal{V} \cong S^1$. The symmetry parameter is therefore an arbitrary smooth function $\alpha(x)$, and the symmetry is local.

When α depends on x, the kinetic term $r_0^2\,\partial_\mu\theta\,\partial^\mu\theta$ is not invariant: under $\theta(x) \mapsto \theta(x) + \alpha(x)$ the derivative picks up $\partial_\mu\alpha(x)$. The geometric resolution is that $\theta(x)$ is a section of a U(1) principal bundle over M, and the correct derivative is the gauge-covariant one $D_\mu\theta = \partial_\mu\theta - A_\mu$, where $A_\mu(x)$ is the connection (the photon field). Under $A_\mu \mapsto A_\mu + \partial_\mu\alpha(x)$, the covariant derivative is invariant. The field strength $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the curvature of the connection. Integrating out short-wavelength fluctuations of θ in Z_{QFT} generates the Maxwell kinetic term at quadratic order in A_μ:

$$S_{\mathrm{Maxwell}} = -\frac{1}{4e^2}\int_M F_{\mu\nu}F^{\mu\nu}\sqrt{-g}\;d^4x, \tag{42}$$

with gauge coupling $e^{-2} \propto r_0^2$. The photon is exactly massless because A_μ transforms inhomogeneously and cannot acquire a gauge-invariant mass term; minimal coupling of charged fields to A_μ via D_μ follows from gauge covariance of the projected kinetic terms. The gauging is not an additional postulate: it is the consequence of Π assigning an independent vacuum phase to each spacetime point.

8.3. $\mathrm{SU}(2)_L$ from the threshold topology: the electroweak doublet

The single complex field ψ is the leading-order description of the modal order parameter. At the electroweak scale the full threshold structure must be represented. The threshold τ is a hypersurface in $\mathcal{C}$; near any crossing point the configuration explores the level set $T(C) = \tau$. The minimal order parameter consistent with both $\mathrm{U}(1)_Y$ and $\mathrm{SU}(2)_L$ is a complex doublet:

$$H = \begin{pmatrix} H^+ \\ H^0 \end{pmatrix} \in \mathbb{C}^2, \qquad |H|^2 = |H^+|^2 + |H^0|^2. \tag{43}$$

The scalar ψ of Section 7 is the neutral component H^0 in the broken phase; H is the Standard Model Higgs doublet.

For $H \in \mathbb{C}^2 \cong \mathbb{R}^4$, the vacuum manifold after crossing is:

$$\mathcal{V}_{\mathrm{EW}} \;=\; \left\{\, H \in \mathbb{C}^2 \;:\; |H|^2 = v^2/2 \,\right\} \;\cong\; S^3, \tag{44}$$

where $v = r_0\sqrt{2} = 246$ GeV is the electroweak vev. The three-sphere S^3 is diffeomorphic to the group manifold of SU(2):

$$S^3 \;\cong\; \mathrm{SU}(2), \tag{45}$$

and its isometry group is $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$. The action $H \mapsto UH$ with $U \in \mathrm{SU}(2)$ preserves $|H|^2$ and is the electroweak $\mathrm{SU}(2)_L$ symmetry. Three of the four real degrees of freedom in H are Goldstone bosons of the breaking $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{em}}$; they are absorbed by the $W^\pm$ and Z bosons, giving them mass. The remaining degree of freedom is the physical Higgs boson of mass $m_h^2 = 4\varepsilon$, identified in equation (38).

The $\mathrm{SU}(2)_R$ isometry is broken by the directional character of the threshold crossing: $\varepsilon = T(C) - \tau > 0$ is an oriented statement in which tension exceeds the threshold from below, selecting a preferred chiral handedness. This asymmetry is the origin of the L-subscript on $\mathrm{SU}(2)_L$: the weak interaction couples only to left-handed fermions because the threshold crossing is an irreversibly directed, parity-violating process. Parity violation in the weak sector is therefore not an external assumption but a theorem of the asymmetric modal threshold.

8.4. $\mathrm{SU}(3)_c$ from substrate dimensionality: colour as compact holonomy

Section 11 derives that the substrate possesses $N = 3$ compact dimensions from the dimensional reduction identity $d_0^N = \ell_*^{N+2}/\ell_P^2$ with $N = 3$, giving $d_0 \approx 7$ nm. The number $N = 3$ is a derived output of PST. It equals the number of QCD colours.

The identification is structural. The N compact substrate dimensions form a fibre over each point of the projected spacetime M. The automorphism group of a compact complex fibre $\mathbb{C}^N$ preserving the Hermitian structure is $\mathrm{U}(N) = \mathrm{U}(1) \times \mathrm{SU}(N)$. For $N = 3$:

$$\mathrm{Aut}\big(\mathbb{C}^3,\, \langle\cdot,\cdot\rangle\big) \;=\; \mathrm{U}(3) \;=\; \mathrm{U}(1) \times \mathrm{SU}(3)_c. \tag{46}$$

The $\mathrm{SU}(3)_c$ factor is the colour gauge group; the U(1) factor contributes to hypercharge. Configurations invariant under $\mathrm{SU}(3)_c$ project to colour singlets (leptons, Higgs, gauge bosons in the adjoint); configurations in the fundamental **3** project to colour triplets (quarks).

Three corollaries follow. *(i)* Quark charges are fractional: a colour-neutral baryon composed of $N_c = 3$ quarks must carry integer electric charge, forcing each quark to carry a charge that is a multiple of $1/3$. Fractional quark charges are a consequence of $N_c = 3$, not a separate postulate. *(ii)* The same integer N governs the Casimir discreteness scale, the colour multiplicity, and the number of spatial dimensions: three observational handles on a single substrate parameter. *(iii)*

Colour confinement at low energies follows from the non-Abelian structure of $\mathrm{SU}(3)_c$, whose running coupling grows in the infrared; PST inherits this without postulating confinement independently.

8.5. The Standard Model gauge group as substrate automorphisms

Assembling the three factors, the automorphism group of the instantiated vacuum is:

$$G_{\mathrm{SM}} \; = \; \mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y, \tag{47}$$

where $\mathrm{U}(1)_Y$ arises from gauging the vacuum-manifold phase rotation, $\mathrm{SU}(2)_L$ from the $S^3 \cong \mathrm{SU}(2)$ topology of the electroweak vacuum manifold, and $\mathrm{SU}(3)_c$ from the $N = 3$ compact substrate fibre. None of the three factors is postulated; each is the automorphism group of a structural object that PST derives independently.

The four fundamental forces are all contained within the emergence chain:

- **Gravity**, the gradient term of $\mathcal{F}$ projected to the Einstein–Hilbert action; $G_{\mu\nu}$ uniquely fixed by diffeomorphism invariance and Lovelock's theorem (Section 7).
- **Electromagnetism**, the $\mathrm{U}(1)_{\mathrm{em}}$ gauge theory from gauging the Goldstone phase mode of $\mathcal{V} \cong S^1 \subset \mathcal{V}_{\mathrm{EW}} \cong S^3$; the photon is the connection on the U(1) bundle over M.
- **The weak force**, the $\mathrm{SU}(2)_L$ gauge theory from the $S^3 \cong \mathrm{SU}(2)$ vacuum manifold; left-chiral because the threshold crossing is directional; $W^{\pm}$ and Z masses from the Higgs mechanism on $\mathcal{V}_{\mathrm{EW}}$.
- **The strong force**, the $\mathrm{SU}(3)_c$ gauge theory from the $N = 3$ compact substrate fibre; eight massless gluons; confinement from the non-Abelian running coupling.

8.6. Gauge anomaly cancellation as a structural theorem

A gauge theory is quantum-mechanically consistent only if its triangle anomalies vanish: the Ward identities of the gauge currents must hold at the quantum level, otherwise the gauge symmetry is destroyed by radiative corrections and the theory is ill-defined. In the Standard Model, anomaly cancellation holds by a precise numerical relationship between the number of colours and the hypercharge assignments of quarks and leptons. In PST, this is a theorem.

Express the Standard Model fermion content of one generation as left-handed Weyl spinors, replacing each right-handed field χ_R by its left-handed conjugate χ_R^c with opposite quantum numbers. With N_c colours, the fields and their hypercharges are: the quark doublet Q_L (N_c colours, $Y = +1/3$), the right-handed up quark $(u_R)^c$ (N_c colours, $Y = -4/3$), the right-handed down quark $(d_R)^c$ (N_c colours, $Y = +2/3$), the lepton doublet L ($Y = -1$), and the right-handed charged lepton $(e_R)^c$ ($Y = +2$). The two independent anomaly conditions that are not automatically satisfied are:

$$\big[\mathrm{U}(1)_Y\big]^3 = 0: \quad N_c\Big[2\big(\tfrac{1}{3}\big)^3 + \big(-\tfrac{4}{3}\big)^3 + \big(\tfrac{2}{3}\big)^3\Big] + 2(-1)^3 + (2)^3 \;=\; -2N_c + 6, \tag{48}$$

$$\big[\mathrm{SU}(2)_L\big]^2\big[\mathrm{U}(1)_Y\big] = 0: \quad N_c \cdot \tfrac{1}{3} + 1 \cdot (-1) \;=\; \tfrac{N_c}{3} - 1. \tag{49}$$

Both vanish if and only if $N_c = 3$. Since PST derives $N_c = N_{\mathrm{compact}} = 3$ in Section 11, the following holds as a theorem:

> *A precausal substrate with N compact dimensions projects to a gauge theory with N colours. The gauge anomaly conditions require $N = 3$ for the projected theory to be quantum-mechanically consistent. PST derives $N = 3$ independently from the Casimir constraint. The Standard Model is anomaly-free because the substrate parameter that fixes $d_0 \approx 7$ nm also selects the unique consistent colour multiplicity.*

Two further conditions are automatically satisfied for any N_c: $[\mathrm{SU}(3)_c]^3 = 0$ because the quark sector contains equal numbers of fundamental and anti-fundamental representations, and $[\mathrm{SU}(2)_L]^3 = 0$ because SU(2) has a vanishing cubic Casimir. The mixed gravitational anomaly $[\mathrm{grav}]^2[\mathrm{U}(1)_Y] \propto \sum_i Y_i = 0$ also vanishes per generation, independently of N_c.

8.7. The Standard Model Lagrangian as leading-order projection

The Standard Model Lagrangian is the leading-order term in the expansion of the instantiated action $S[\psi, g, A^a_\mu]$ around the vacuum $\mathcal{V}_{\mathrm{EW}}$:

$$\mathcal{L}_{\mathrm{SM}} \;=\; -\tfrac{1}{4}F^a_{\mu\nu}F^{a\,\mu\nu} + \bar{\psi}\, i\, \not{D}\, \psi + |D_\mu H|^2 - \mathcal{F}(H) + \mathcal{L}_{\mathrm{Yukawa}}, \tag{50}$$

where $F^a_{\mu\nu}$ runs over all gauge factors of G_{SM}, $\mathcal{F}(H)$ is the modal potential (33) evaluated on the Higgs doublet, and D_μ is the G_{SM}-covariant derivative. The gauge kinetic terms arise from the Maxwell construction applied to each factor; the scalar potential is the modal potential extended to $H \in \mathbb{C}^2$; the Higgs kinetic term $|D_\mu H|^2$ encodes minimal coupling to all gauge fields. Corrections from higher substrate-expansion terms are suppressed by $(kd_0)^2$ and contribute observably only in the Casimir regime near $d_0 \approx 7$ nm. The Yukawa sector $\mathcal{L}_{\mathrm{Yukawa}} = y_{ij}\bar{\psi}_i H \psi_j + \mathrm{h.c.}$ and the fermionic kinetic term require the fermionic sector.

8.8. The fermionic sector: structure and the sharpened open problem

The excitations of the complex doublet $H \in \mathbb{C}^2$ are integer-spin fields: the Higgs scalar, the gauge bosons, and the graviton. Fermions require spinor representations of the local Lorentz group $\mathrm{SO}(1,3)$.

The spinor bundle exists. Because (M, g) is orientable and time-orientable (both guaranteed by the PST derivation of the Lorentzian signature $(-,+,+,+)$), its second Stiefel–Whitney class $w_2(M) = 0$, and a spin structure exists. The vierbein $e^a{}_\mu$ satisfying $g_{\mu\nu} = e^a{}_\mu e^b{}_\nu \eta_{ab}$ is a derived quantity (the square root of the projected metric), not an independent primitive. The

v9.9

spinor bundle $\mathcal{S} = \mathcal{S}^+ \oplus \mathcal{S}^-$ on (M, g) is therefore structurally available. Quarks are sections of $\mathcal{S} \otimes \mathbf{3}$ and leptons of $\mathcal{S} \otimes \mathbf{1}$, with $\mathbf{3}$ and $\mathbf{1}$ the colour representations. The Dirac operator on $\mathcal{S}$ formed from the projected metric generates the Dirac equation at the kinematic level (the equation's operator structure follows from the spinor bundle and the projected metric); the dynamical question of whether the substrate produces fermionic excitations that populate this equation is a separate and open problem.

The mechanism for fermionic excitations from the substrate is open. The order parameter H is bosonic; its excitations obey Bose statistics. Grassmann (anti-commuting) integration variables in Z_{QFT} are required for half-integer spin. The threshold vacuum manifold $\mathcal{V}_{\mathrm{EW}} \cong S^3 \cong \mathrm{SU}(2)$ admits half-integer representations; a spin structure on S^3 exists without topological obstruction ($\pi_1(S^3) = 0$). The precise open question is whether the Bernoulli measure μ gives non-zero weight to path-integral sectors transforming in the spinor representation $j = 1/2$ of the threshold SU(2), and whether Π projects those sectors to fermionic fields satisfying the Dirac equation on M. An affirmative answer would make the number of generations N_{gen}, the Yukawa hierarchy, and the CKM matrix the next layer of open problems. These are listed in Section 14.

8.9. Layman's Summary

This section bridges PST to the particle physics of the Standard Model. It shows how quantum field theory's path integral arises from the substrate's partition function, and derives the four fundamental forces from the geometry of the substrate: electromagnetism from the circular symmetry of the vacuum phase, the weak force from the three-sphere topology of the electroweak vacuum, and the strong force from the three compact extra dimensions of the substrate. Crucially, a consistency requirement called gauge anomaly cancellation (which demands that the number of quark colour charges be exactly three) is not an accident but a theorem: the same $N = 3$ that fixes the Casimir correction also makes the quantum theory self-consistent.

9. Physical Predictions

A foundational theory makes predictions at two levels. Structural predictions follow from the form of the theory alone, independently of the specific content of any particular precausal configuration C. Specific predictions require knowledge of C and concern the particular values of physical parameters in our universe. This section focuses on structural predictions: consequences that hold for any instantiated geometry, regardless of the detailed tension distribution that produced it. These are the predictions that a future derivation of the Standard Model from PST would inherit as constraints; they are also the predictions through which PST can in principle be distinguished from competing frameworks experimentally.

A second point about the character of predictions is worth stating explicitly. PST is not primarily a replacement for existing physical theories at the level at which those theories are predictive. General Relativity and Quantum Mechanics already give correct answers to the questions they ask. PST operates at the level beneath those questions, at the level of why those theories work, why the spacetime arena exists, why matter inhabits it, and why the symmetries they assume hold. The predictions in this section are therefore predictions about the domain where existing theories break down, approach their limits, or require assumptions that PST renders derivable.

9.1. Quantum Fluctuations

Near the modal threshold, where $\varepsilon(C) = T(C) - \tau \to 0$, the order parameter ψ is in the critical region: the two minima of the modal potential are about to separate but have not yet done so. The potential near $\psi = 0$ is nearly flat in the direction of bifurcation, meaning that the order parameter requires almost no energy to shift significantly. Small fluctuations in the tension δT produce large fluctuations in the instantiated metric:

$$\delta g_{\mu\nu} \;\sim\; \eta_{\mu\nu} \cdot \left. \frac{\delta T}{2\,|T - \tau_{\text{null}}|^{1/2}} \right|_{T \to \tau_{\text{null}}} \tag{51}$$

This is the PST interpretation of quantum fluctuations. They are not a fundamental feature of the ontology, not irreducible randomness built into the fabric of reality. They are the geometric signature of tension configurations that are close to the modal threshold: configurations whose instantiation is marginal, whose order parameter has not settled firmly into one of the two degenerate minima.

The Heisenberg uncertainty principle, $\Delta x\, \Delta p \geq \hbar/2$, acquires a natural interpretation in this picture. Position is registered by the realisation map ρ, which assigns to each property a point in the instantiated manifold. Momentum is the projection of the directional tension gradient via Π. Near the threshold, ρ is sensitive to small variations in $T(C)$ because the denominator in equation (18) is small near τ_{null}. This sensitivity means that determining position precisely, which amounts to resolving ρ to high accuracy, requires probing configurations at separations close to τ_{null}, which in turn amplifies the uncertainty in the tension gradient that projects as momentum. The uncertainty relation is a lower bound on the product of these two projection sensitivities; it is a structural property of Π near the threshold, not an empirical law imposed from outside.

Wave-particle duality has the same precausal interpretation. A configuration that is deep in the instantiated regime, far from the threshold with $\varepsilon(C) \gg 0$, has a sharply defined order parameter and projects through Π as a localised density: a particle. A configuration hovering near the threshold has a diffuse, oscillating order parameter and projects as a spatially extended metric fluctuation: a wave. The transition between wave behaviour and particle behaviour is not a collapse caused by measurement; it is the settling of the order parameter from the near-threshold critical region into one of the two stable minima. Measurement in PST is not a physical interaction that disturbs a pre-existing state; it is the process by which a near-threshold configuration is brought sufficiently far from τ that its projection becomes localised.

v9.9

9.2. Spacetime Curvature

At macroscopic scales, where $\varepsilon(C) \gg 0$ and the order parameter has settled firmly into the instantiated regime, the dominant effect of the tension distribution is not fluctuation but curvature. A region of the instantiated manifold is curved to the extent that the tension function $T(C)$ varies nonuniformly across the precausal configurations that project into it. From the tension-to-metric map (18) and the definition of the Ricci scalar:

$$R\big(\Phi(C)\big) \;\propto\; \Pi\big(\nabla^2_{\mathcal{C}} T(C)\big) \tag{52}$$

Curvature is the projection of the configuration-space Laplacian of asymmetric tension onto the instantiated manifold. The Laplacian $\nabla^2_{\mathcal{C}}$ acts on the discrete power set $\mathcal{P}(D)$; its continuum realisation via Π is an open problem listed in Section 14. A perfectly uniform tension distribution, in which $T(C)$ is the same for all configurations, projects as flat Minkowski space. Every departure from flatness in the instantiated geometry corresponds to a nonuniformity in $T(C)$. Massive objects are therefore not the source of curvature in a causal sense; they are the instantiated projection of regions where the precausal tension distribution was particularly nonuniform, and the curvature they produce is simply the geometric expression of that nonuniformity.

Gravitational waves follow immediately from this picture. A propagating disturbance in $T(C)$ in the precausal configuration space projects as a propagating disturbance in $g_{\mu\nu}$ in the instantiated manifold. Gravitational waves are not ripples in a pre-existing fabric; they are oscillating tension gradients in configuration space manifesting as metric perturbations. Their propagation at the speed of light follows from the fact that the null threshold τ_{null} is an isotropic property of the projection: disturbances propagate at the speed set by the boundary between timelike and spacelike separations, which is fixed by τ_{null} alone. The recent detection of gravitational waves by LIGO [46] is fully consistent with this picture; the waves are a direct experimental observation of propagating tension gradients in the precausal substrate.

Black holes are the extreme limit of this gradient picture: they correspond to configurations in which $\nabla^2_{\mathcal{C}} T(C) \to \infty$ at a point in configuration space. The event horizon is the projection of the set of configurations at which the tension gradient reaches the null threshold condition, beyond which no precausal configuration can project outward. The Hawking temperature of a black hole acquires a precausal interpretation as the thermal fluctuation of near-threshold configurations at the horizon, consistent with the thermodynamic derivation of Jacobson [18] but now grounded in the precausal structure of Π rather than in a phenomenological entropy argument.

9.3. The weak-field limit and the recovery of Newtonian gravity

Equation (28) is the Einstein field equation. For the claim that PST derives GR to be substantive, rather than merely renaming GR's content, the field equation must reduce to Newtonian gravity in the non-relativistic, weak-field limit. The reduction is standard in GR; it is reproduced here in PST notation to confirm that no additional structure interferes.

In the weak-field, slow-motion regime, write $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$ with $|h_{\mu\nu}| \ll 1$. The projected tension reduces to a non-relativistic matter density: $\Pi(T(C))(x) \approx \rho(x)$, so that $T_{00} = \rho$ and all spatial components are negligible. Working in harmonic gauge with $\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h$, the linearised field equations become:

$$\nabla^2 \bar{h}_{00} \;=\; -16\pi G\,\rho \tag{53}$$

Setting $\bar{h}_{00} = -4\phi$, where ϕ is the Newtonian gravitational potential ($g_{00} \approx -(1 + 2\phi)$), equation (53) reduces to:

$$\nabla^2 \phi \;=\; 4\pi G\,\rho \tag{54}$$

Poisson's equation for Newtonian gravity is recovered exactly, with G playing the role of the unit-conversion factor between the projected matter density $\rho = \Pi(T(C))$ and the metric perturbation ϕ. At the next post-Newtonian order the same field equations reproduce perihelion precession ($\gamma = \beta = 1$) and gravitational light bending at the value GR predicts, since equation (28) is the Einstein equation; PST introduces no additional fields or couplings that would shift the post-Newtonian parameters away from their GR values. The sole physical differences from standard GR appear at scales near d_0 (the Casimir correction of Section 10) and in the ontological provenance of G and Λ.

9.4. Stress-Energy as Projected Tension

The unification of geometry and matter as coprojections of the same precausal tension resolves a longstanding conceptual problem in theoretical physics. In General Relativity, the statement "matter curves spacetime" is taken as a foundational description: matter and spacetime are separate ontological kinds, and one acts on the other. In Quantum Field Theory, fields and their excitations are defined on a background spacetime, which is treated as a fixed classical arena. The two descriptions are mutually inconsistent at a deep level: the quantum fields of QFT back-react on the spacetime metric, but the resulting semiclassical equations are not derivable from a unified action, and the renormalisation of the stress-energy tensor introduces ambiguities that have no physical resolution within the framework.

PST dissolves this problem by removing the distinction between the two categories. Geometry and stress-energy are coprojections: both are outputs of Φ and Π acting on the same input $T(C)$. The phrase "matter curves spacetime" is a useful shorthand within the instantiated domain, but its deeper meaning is that the same precausal tension structure that projects as matter also projects as curvature, at every point simultaneously. There is no back-reaction because there are not two things acting on each other; there is one projection producing two aspects of a single geometric object.

Mass, in this picture, is a measure of how far a configuration sits from the modal threshold. A configuration deep in the instantiated regime, with $\varepsilon(C) \gg 0$, has a large order parameter $r_0 \propto \sqrt{\varepsilon}$ and projects as a dense, heavy field configuration. A configuration hovering near the threshold has a small order parameter and projects as a light or massless excitation. Massless fields, including the photon, correspond to configurations that project through Π at exactly the null threshold τ_{null}: they occupy the boundary between timelike and spacelike separation, which is precisely the condition for propagation at the speed of light. The masslessness of light is not a postulate; it is the projection of the null threshold condition into the instantiated geometry.

9.5. Vacuum Phenomena and the Casimir Effect

Standard quantum field theory predicts a Casimir pressure between two parallel perfectly conducting plates at separation d of [11]:

$$P_{\text{standard}}(d) = -\frac{\pi^2 \hbar c}{240\, d^4} \tag{55}$$

At $d = 100$ nm this evaluates to approximately -13 Pa, measurable experimentally [12, 13]. The effect has been measured to sub-percent accuracy at separations 160–750 nm [14, 15].

PST provides a precausal account of this effect. The conducting plates impose boundary conditions on the projection operator Π, restricting which sub-configurations C' contribute to the kernel integral at points x between the plates. In standard QFT, equivalent boundary conditions are imposed on field modes, and the two descriptions agree at leading order. The departure arises because the kernel $K(x, C') = \delta(x - \rho(C'))$ has a spatial resolution structure set by the characteristic scale d_0 at which the realization map $\rho : D \to M$ becomes sensitive to boundary geometry. At separations $d \gg d_0$ the two descriptions are equivalent; at $d \sim d_0$ they diverge.

By dimensional analysis, and the parity constraint that the correction must be even in d_0/d (odd powers would require a preferred orientation inconsistent with the isotropy of the precausal substrate), the leading PST correction to the Casimir pressure is:

$$P_{\text{PST}}(d) = P_{\text{standard}}(d) \left[1 + \xi \left(\frac{d_0}{d} \right)^2 + \mathcal{O}\left(\left(\frac{d_0}{d} \right)^4 \right) \right] \tag{56}$$

where ξ is a dimensionless coefficient of order unity set by the structure of K. The correction term scales as d^{-6}:

$$\delta P_{\text{PST}}(d) = -\frac{\xi\, \pi^2 \hbar c\, d_0^2}{240\, d^6} \tag{57}$$

This d^{-6} scaling is the distinctive experimental signature of PST.

9.6. The parameter-free exponent and the structure of falsifiability

A careful reader will observe that equation (57) contains two quantities that PST does not derive from first principles at this stage: the coefficient ξ and the scale d_0. It is important to be precise about what this means for the theory's falsifiability, because the situation is more favourable than it might appear.

The correction $\delta P_{\mathrm{PST}} \propto \xi\, d_0^2\, d^{-6}$ involves two parameters that determine its *amplitude*: the dimensionless coefficient ξ and the discreteness scale d_0. Neither is freely adjustable in the sense that could hide a failure: d_0 is bounded above by existing Casimir data (equation (58)) and computed from the dimensional reduction mechanism in Section 10; ξ is of order unity by a structural argument (it is set by the kernel K of the substrate's configuration-space geometry) and cannot be made small without additional theoretical input that PST does not have. Nevertheless, it is true that neither value is yet computed from first principles, and one should be honest that the *amplitude* of the prediction therefore carries residual parametric uncertainty.

What does *not* depend on ξ or d_0 is the power-law *exponent*. The exponent -6 in $\delta P_{\mathrm{PST}} \propto d^{-6}$ follows from dimensional analysis and the parity constraint alone: the correction must be even in d_0/d (odd powers would require a preferred orientation of the substrate, inconsistent with the isotropy established in Section 3), and the leading even power is $(d_0/d)^2$, which, multiplied by the standard d^{-4} Casimir scaling, gives d^{-6}. This argument uses only the existence of a length scale d_0 and the symmetry of the substrate; it does not use the values of ξ or d_0. The exponent -6 is therefore a *categorical* prediction: it cannot be changed by adjusting ξ upward or downward, by moving d_0 within its allowed range, or by any other parametric variation within the theory.

The experimental consequence is that the test of PST's Casimir prediction has two logically distinct layers. The *first layer* is qualitative and parameter-free: does the measured Casimir pressure, after subtracting the standard d^{-4} term and all known corrections (finite conductivity, surface roughness, thermal fluctuations), show a residual that scales as d^{-6}? If the residual scales as d^{-4}, or as d^{-5}, or follows any other power law, PST is falsified regardless of any choice of ξ or d_0. This is a categorical falsification: the theory predicts a specific exponent and the measurement returns a different one. The *second layer*, conditioned on the first being passed, is quantitative: does the amplitude of the d^{-6} correction match the prediction $\xi\, d_0^2$ with values of $\xi \sim 1$ and $d_0 \approx 7$ nm? Failure at this layer would constrain ξ and d_0 but would not immediately falsify the exponent structure.

The ratio test of equation (59) is designed to probe the first layer cleanly. By measuring the ratio $P(d_1)/P(d_2)$ at two separations d_1, d_2 where both the standard Casimir term and the PST correction are non-negligible, one can extract the scaling exponent from the data directly, without needing to know the absolute amplitude. If the ratio test is consistent with d^{-6} scaling, the second layer then constrains ξ and d_0. If it is inconsistent, PST is falsified at the exponent level.

For comparison: the gravitational wave tests of General Relativity similarly separate a categorical prediction (the waveform morphology depends on a specific combination of mass and spin, not on any free parameter) from a quantitative one (the amplitudes constrain the source parameters). A theory that predicts a specific power-law exponent as a structural consequence is making a categorical claim about the form of its correction, not merely about its magnitude. PST's claim that the Casimir correction scales as d^{-6} is of this categorical kind.

The d^{-6} scaling is separable from the standard d^{-4} Casimir dependence by fitting measurements at multiple separations, and it is distinguishable from all other known corrections: finite conductivity corrections scale as d^{-4} with a temperature- and material-dependent prefactor; surface roughness introduces corrections proportional to the power spectrum of the surface at the relevant spatial frequency; thermal (Lifshitz) corrections have a distinct temperature dependence that can be isolated by varying the plate temperature at fixed separation. Each competing correction has a known functional form and known parametric dependences, so a d^{-6} signal cannot be mimicked by their combination without an extremely implausible fine-tuned cancellation [15].

9.7. Observational constraint

Precision Casimir measurements at separations 160–750 nm report approximately 1% agreement with standard QFT [14, 15]. Requiring $|\xi|(d_0/d)^2 < 0.01$ at $d = 160$ nm gives:

$$d_0 < \frac{16\ \text{nm}}{\sqrt{|\xi|}} \tag{58}$$

For $|\xi| \sim 1$ this constrains $d_0 < 16$ nm. The PST correction becomes detectable at the 1% level for separations $d \lesssim 10\, d_0$, placing the accessible signal at $d \lesssim 160$ nm at the precision of existing measurements. A ratio test between measurements at two separations d_1 and d_2 provides a clean, apparatus-independent signature:

$$\frac{P(d_1)/P_{\text{standard}}(d_1)}{P(d_2)/P_{\text{standard}}(d_2)} - 1 \propto \left(\frac{d_0}{d_1}\right)^2 - \left(\frac{d_0}{d_2}\right)^2 \tag{59}$$

This represents a falsifiable prediction distinguishable from both standard QFT and from Lifshitz-theory corrections.

The value of d_0 is not a free parameter. Section 10 derives it from PST's dimensional reduction mechanism, the compactification of three substrate dimensions during modal sublimation, and shows that $d_0 \approx 7$ nm when the fundamental modal scale is identified with the electroweak condensate. At this value, the correction at $d = 160$ nm is $|\xi| \times 0.19\%$, below current precision; at $d = 50$ nm it reaches $|\xi| \times 2\%$, within reach of next-generation experiments; and at $d = 20$ nm it reaches $|\xi| \times 12\%$, unambiguous against all known competing corrections. The Casimir measurement is accordingly a genuine predicted detectable signal, not merely a formal bound.

9.8. Layman's Summary

Having established the theoretical structure, this section identifies what PST predicts that can actually be measured. General Relativity and quantum field theory already predict gravitational curvature and quantum fluctuations correctly; PST recovers both. Where PST departs from standard physics is at very short distances: because the substrate has a finite grain size, the vacuum is not perfectly smooth, and this deviation appears in the Casimir effect, the tiny attractive force between two uncharged metal plates held very close together. PST predicts that at separations below about 100 nanometres the force scales differently from standard predictions, by an amount large enough to detect with near-future precision instruments.

10. Grounding d_0: Modal Coherence, Dimensional Reduction, and the Casimir Prediction

The Casimir correction derived in Section 9 has a distinctive d^{-6} scaling whose amplitude is proportional to d_0^2. Section 7 related G and d_0 via $G = \alpha_G d_0^2$ in natural units, but did not constrain α_G. A reader might assume $\alpha_G \sim 1$, which would place $d_0 \approx \ell_P \approx 1.6 \times 10^{-35}$ m. At $d = 160$ nm this gives $(d_0/d)^2 \sim 10^{-56}$: the correction would be fifty-six orders of magnitude below the leading term and permanently undetectable.

The assumption $\alpha_G \sim 1$ is not a result of PST. It was never derived; it was imported as a prior that the only relevant length scale in the theory is ℓ_P. This section shows the prior is wrong, and by how much: the correct value of α_G, derived from PST's own dimensional reduction mechanism, is approximately 10^{-54}, placing d_0 in the nanometre range twenty-six orders of magnitude above ℓ_P. At this value the Casimir correction is a genuine predicted detectable signal, not a formal statement about an unobservably small parameter.

10.1. Two distinct scales: microscopic grain versus coherence length

The conflation of d_0 with ℓ_P rests on identifying two physically distinct lengths: the *microscopic grain* of the substrate and the *coherence length* of the modal condensate. The closest condensed-matter analogue makes the distinction precise.

In BCS superconductivity the ionic crystal has lattice constant $a \approx 0.3$ nm. This is the microscopic scale of the problem. The characteristic scale of the superconducting order, however, is the Cooper pair coherence length:

$$\xi_{\mathrm{BCS}} = \frac{\hbar v_F}{\pi \Delta} \tag{60}$$

where v_F is the Fermi velocity and Δ the superconducting gap. Because Δ/E_F is exponentially small for conventional superconductors ($\sim 10^{-3}$–10^{-5}), the coherence length satisfies $\xi_{\mathrm{BCS}} \sim 10^3 a$ to $10^6 a$. The scale at which the condensate is probed exceeds the lattice constant by three to six orders of magnitude. This is not an accident: it is the universal consequence of near-criticality. Near a second-order phase transition the correlation length diverges; a system barely below its critical point has a coherence length that can be macroscopically large even when the underlying lattice is atomic.

PST's modal condensate has the same architecture. The substrate carries a microscopic scale ℓ_*, the intrinsic grain of the distinction space D at which primitive differences are registered. Nothing in the PST axioms forces $\ell_* = \ell_P$; and even if it did, d_0 would not equal ℓ_*. The scale d_0 is defined by the projection operator Π as the finest resolution at which the realisation map $\rho : D \to M$ can distinguish spacetime points. This is the spatial scale over which the modal order parameter ψ varies: its coherence length. In the Landau-Ginzburg functional of Section 3:

$$F[\psi] \;=\; \int \mathrm{d}^4x \left[\tfrac{1}{2}\, c_{LG}\, |\nabla\psi|^2 \;+\; a_{LG}\, |\psi|^2 \;+\; b_{LG}\, |\psi|^4\right] \tag{61}$$

the coherence length of ψ is:

$$d_0 \;=\; \xi_\psi \;=\; \sqrt{\frac{c_{LG}}{2\,|a_{LG}|}} \tag{62}$$

Near the modal threshold, $a_{LG} \propto (T/\tau - 1)$, so $\xi_\psi \to \infty$ as $T \to \tau^+$: the correlation length diverges at the transition, exactly as in any second-order phase transition. The universe exists because $T > \tau$, but nothing requires T to lie far above τ. A substrate that barely exceeds the modal threshold has $d_0 = \xi_\psi \gg \ell_*$. How much larger d_0 is depends on how close to threshold the substrate sits, and this is what the dimensional reduction mechanism determines.

10.2. Modal sublimation as dimensional reduction

The functor $\Phi : \mathcal{S} \to \mathcal{G}$ maps the precausal substrate to an instantiated geometry with four *large* dimensions (three spatial, one temporal), established by the coinstantiation argument of Section 5 and the orbital stability requirement of Section 6. But the substrate configuration space has more than four degrees of freedom.

The tension structure $T : D \times D \to \mathbb{R}$, beyond the scalar order parameter ψ captured by the Landau-Ginzburg functional, carries three independent directional modes corresponding to the three spatial directions of asymmetric tension redistribution. These three modes transform as a vector under the SO(3) spatial symmetry of the instantiated geometry. They do not participate in the macroscopic sublimation that produces the four large spacetime dimensions; instead, they remain compact at the characteristic scale d_0 of their own coherence length.

Modal sublimation therefore produces $(4 + N)$-dimensional geometry with $N = 3$ compact spatial dimensions at scale d_0. The number $N = 3$ is not a free parameter: it equals the dimensionality of the vector representation of the spatial symmetry group SO(3), which is itself fixed by the 3+1 dimensionality established in Sections 5–6. Three large spatial dimensions require exactly three directional tension modes, which compactify as three extra dimensions.

The relationship between the observed four-dimensional Newton's constant G_4 and the $(4+N)$-dimensional fundamental coupling G_* follows from integrating out the N compact dimensions [16]:

$$G_4 \;=\; \frac{G_*}{d_0^N} \qquad \Longleftrightarrow \qquad \ell_P^2 \;=\; \frac{\ell_*^{N+2}}{d_0^N} \tag{63}$$

where $\ell_* \equiv (G_*)^{1/(N+2)}$ in natural units $\hbar = c = 1$. This is the PST dimensional reduction identity. It relates three independently accessible quantities: the Planck length ℓ_P (fixed by measurement of G, $\hbar$, c), the fundamental substrate length ℓ_* (determined below), and the compactification radius d_0 (the prediction, the same d_0 that enters the Casimir correction).

Equation (63) replaces the provisional link between d_0 and ℓ_P established in Section 7. The

v9.9

two relations are consistent:

$$G_4 = \alpha_G \cdot d_0^2 \qquad \text{(Section 7: defining relation)} \tag{64}$$

$$G_4 = \ell_*^{N+2}/d_0^N \qquad \text{(Section 10: dimensional reduction)} \tag{65}$$

Comparing them with $N = 3$ gives $\alpha_G = \ell_*^5/d_0^5 = (\ell_*/d_0)^5$. This confirms that $\alpha_G \ll 1$ whenever $d_0 \gg \ell_*$, which is precisely what dimensional reduction predicts. The provisional assumption $\alpha_G \sim 1$ was equivalent to assuming $d_0 = \ell_*$, i.e., that the compactification scale equals the fundamental substrate grain: a statement that has no justification within the theory.

For $N = 3$, equation (63) yields:

$$d_0^3 = \frac{\ell_*^5}{\ell_P^2} \tag{66}$$

This is the master prediction equation: given ℓ_*, it determines d_0 completely.

A sharper observation is worth stating here. Three equations are now present simultaneously in this paper:

$$d_0 = \sqrt{\frac{c_{LG}}{2\,|a_{LG}|}} \tag{A}$$

$$d_0^3 = \frac{\ell_*^5}{\ell_P^2} \tag{B}$$

$$c_{LG}\, r_0^2\, \ell_*^3 \sim \frac{M_P^2}{16\pi} \tag{C}$$

Equation (A) is the Landau-Ginzburg coherence length (Section 10, equation (62)); equation (B) is the dimensional reduction identity just derived; equation (C) is the gravitational coupling constraint from Section 7 (the projection of the gradient term to the Einstein-Hilbert action, with the volume ℓ_*^3 of the three compact dimensions explicitly included). Together with the Higgs identification $r_0 = v$, $|a_{LG}| = m_h^2/4$, all three equations are expressed in terms of known measured quantities (v, m_h, M_P) and two unknowns (c_{LG} and ℓ_*). Equations (A) and (C) together eliminate c_{LG}, leaving a single equation in ℓ_* alone, which equation (B) then closes.

Equations (A) and (B) alone are sufficient to eliminate d_0 and express ℓ_* as a closed-form function of c_{LG}, m_h, and ℓ_P; that solution is carried out in the subsection "Simultaneous solution" below. Equation (C) then constrains c_{LG} up to the normalization factor of the compact fibre projection integral, and that integral is what Problem 1 of Section 14 asks for.

10.3. The fundamental modal scale and the electroweak identification

In PST, ℓ_* is the length associated with the energy at which the modal condensate forms, the scale at which the order parameter ψ first acquires a non-zero expectation value. The current section uses a structural identification to pin ℓ_* to a known energy scale pending the simultaneous solution described above. What PST can establish without that solution is the structural argument that pins ℓ_* to a known energy scale.

The PST order parameter ψ and the Standard Model Higgs field ϕ occupy the same mathematical role. Both are complex scalars that:

1. inhabit a degenerate minimum of a U(1)-symmetric quartic potential;
2. break that symmetry by acquiring a non-zero vacuum expectation value;
3. produce a massive longitudinal (Higgs/radial) mode and a massless Goldstone direction (photon/S^1 vacuum manifold).

In PST's ontological framework, the Higgs condensate is the four-dimensional projection of the modal condensate: $\phi = \Pi(\psi)$, the order parameter as seen from inside the instantiated geometry. If this identification holds, then the modal condensate and the Higgs condensate share the same vacuum structure, and $M_* = \hbar c/\ell_*$ is the energy scale at which the full precausal modal structure becomes manifest, above the electroweak vev $v = 246$ GeV, where the condensate itself forms, but within the regime constrained by LHC electroweak searches.

A reader will note that the identification now places three distinct energy scales near the condensate, and their consistency should be made explicit. The radial-mode mass derived in Section 7, $m_h^2 = 4\varepsilon$, and the vacuum radius $r_0^2 = \varepsilon/(2b)$ of equation (34) together give $m_h^2 = 8b\,r_0^2$, so the three quantities are not independent: fixing any two fixes the third through the quartic coefficient b. Identifying r_0 with the Higgs vacuum expectation value $v = 246$ GeV and b with one quarter of the Standard-Model Higgs self-coupling, $b = \lambda/4$ with $\lambda \approx 0.13$, returns $m_h = \sqrt{2\lambda}\,v \approx 125$ GeV, the measured Higgs mass, with no additional input. The modal scale $M_* \approx 1.25$ TeV is a fourth and separate quantity: it is not the vev and not the radial-mode mass but the threshold energy $\hbar c/\ell_*$ that sets the substrate grain entering the dimensional-reduction identity (66), and it lies above the condensate observables rather than coinciding with either. The Higgs mass and the vev are the low-energy observables of the modal condensate; M_* is the scale at which the full modal structure, the "new physics" of the condensate's precausal origin, becomes manifest. The three are mutually consistent through $b = \lambda/4$ and $v < M_*$, and only the single ratio M_*/v remains a genuine prediction awaiting the unconditional derivation of ℓ_* (Problem 1).

The Higgs field acquires its vacuum expectation value at $v = 246$ GeV. Precision electroweak data and LHC direct searches constrain new physics associated with the electroweak condensate to begin above approximately 1 TeV.

The choice of M_* requires careful statement. The fundamental modal scale is $M_* = \hbar c/\ell_*$, where ℓ_* is the substrate grain length. A derivation of ℓ_* from first principles within PST has not yet been carried out; it is Open Problem 1 of Section 14. The value $M_* = 1.25$ TeV used in what follows is therefore a *lower bound* obtained from experimental constraints, not a theoretical prediction: LHC direct searches and electroweak precision observables exclude new resonances associated with the condensate below ~ 1 TeV, and 1.25 TeV is chosen as a representative point consistent with those exclusions.

The implication for the Casimir prediction must be stated honestly. The numerical value

v9.9

$d_0 \approx 7$ nm (derived in the subsection below) is conditional on this choice of M_*. Different values of M_* within the LHC-allowed range give different values of d_0 through the dimensional-reduction identity (66); the amplitude of the Casimir correction, which scales as d_0^2, is therefore parametrically sensitive to where M_* sits within its allowed range. A reader who asks "what if $M_* = 2$ TeV?" is asking a legitimate question, and the answer is that d_0 would change and the predicted onset separation would shift accordingly.

What does not change is the power-law exponent. As established in Section 9, the d^{-6} scaling of the PST Casimir correction is determined by the parity constraint and dimensional counting, not by the value of M_*. The exponent -6 is therefore a prediction that is robust across the entire LHC-allowed range of M_*; only the amplitude and onset separation depend on M_*. The experimental test that directly targets the exponent, fitting the separation dependence of the Casimir pressure residual across multiple plate distances, is therefore insensitive to the choice of M_*, and it is the appropriate primary test of PST's Casimir prediction.

A further constraint tightens the allowed range of M_*. Once ℓ_* is derived from PST's first principles (Problem 1), the ratio d_0/ℓ_* will be fully determined by the dimensional-reduction identity (66), and d_0 will be a derived number with no residual parametric freedom. The Casimir amplitude will then become a genuine second parameter-free prediction, supplementing the exponent. The present situation, in which the amplitude is conditional on M_* while the exponent is not, is an intermediate state of the theory; Problem 1 is the step that closes the remaining freedom.

Taking $M_* = 1.25$ TeV gives:

$$\ell_* \;=\; \frac{\hbar c}{M_*} \;=\; \frac{197.3\ \text{MeV}\cdot\text{fm}}{1250\ \text{GeV}} \;\approx\; 1.578 \times 10^{-19}\ \text{m} \tag{67}$$

10.4. Simultaneous solution of equations (A) and (B): the closed-form expression for M_*

With the Higgs identification $|a_{LG}| = m_h^2/4$ in place, equations (A) and (B) can be solved simultaneously for ℓ_*, eliminating d_0 algebraically.

Step 1: substitute the Higgs identification into equation (A).

$$d_0^2 \;=\; \frac{c_{LG}}{2\,|a_{LG}|} \;=\; \frac{2\,c_{LG}}{m_h^2} \tag{68}$$

Step 2: eliminate d_0. Raising equation (68) to the third power and equation (B) to the second power both produce d_0^6:

$$d_0^6 = \left(\frac{2\,c_{LG}}{m_h^2}\right)^3 = \frac{8\,c_{LG}^3}{m_h^6} \tag{69}$$

$$d_0^6 = \left(\frac{\ell_*^5}{\ell_P^2}\right)^2 = \frac{\ell_*^{10}}{\ell_P^4} \tag{70}$$

Setting (69) equal to (70):

$$\ell_*^{10} = \frac{8\,c_{LG}^3\,\ell_P^4}{m_h^6} \tag{71}$$

Step 3: solve for ℓ_* and M_*. Taking the tenth root:

$$\ell_* = \frac{(2\,c_{LG})^{3/10}\,\ell_P^{2/5}}{m_h^{3/5}} \tag{72}$$

Converting to an energy scale $M_* = \hbar c/\ell_*$ and writing $M_P = \hbar c/\ell_P$:

$$\boxed{M_* = m_h^{3/5}\,M_P^{2/5}\,(2\,c_{LG})^{-3/10}} \tag{73}$$

This is the closed-form result. The fundamental modal scale M_* is expressed entirely in terms of the measured Higgs mass m_h, the measured Planck mass M_P, and the gradient coefficient c_{LG} of the modal condensate. Note that c_{LG} is dimensionless: from equation (68), $c_{LG} = m_h^2 d_0^2/2$, a ratio of squared mass scales, and the right-hand side of equation (73) is accordingly dimensionally consistent, $[\mathrm{GeV}^{3/5} \cdot \mathrm{GeV}^{2/5} \cdot 1] = \mathrm{GeV}$.

Inverting equation (73) gives the constraint that c_{LG} must satisfy to reproduce a given M_*:

$$c_{LG} = \frac{m_h^2\,M_P^{4/3}}{2\,M_*^{10/3}} \tag{74}$$

Using the LHC lower bound $M_* = 1.25$ TeV as the fiducial value:

$$m_h^2 = (125.1\ \mathrm{GeV})^2 \approx 1.565 \times 10^4\ \mathrm{GeV}^2$$

$$M_P^{4/3} = (1.221 \times 10^{19}\ \mathrm{GeV})^{4/3} \approx 2.81 \times 10^{25}\ \mathrm{GeV}^{4/3}$$

$$M_*^{10/3} = (1.25 \times 10^3\ \mathrm{GeV})^{10/3} \approx 2.10 \times 10^{10}\ \mathrm{GeV}^{10/3}$$

$$c_{LG} \approx \frac{1.565 \times 10^4 \times 2.81 \times 10^{25}}{2 \times 2.10 \times 10^{10}} \approx 1.0 \times 10^{19} \tag{75}$$

This is a dimensionless number. Its magnitude reflects the ratio of the compactification scale to the electroweak inverse mass: $c_{LG} = m_h^2 d_0^2/2 \approx (m_h d_0)^2/2$, and $m_h d_0 \approx 4.4 \times 10^9$ (the compactification scale $d_0 \approx 7$ nm spans approximately 4.4×10^9 Higgs-Compton wavelengths $1/m_h$). This large but structural number encodes the same hierarchy that makes gravity weak: the weakness of G and the value of c_{LG} are two faces of the same ratio d_0/ℓ_*.

For M_* ranging over the LHC-allowed interval 1.0–1.5 TeV, the corresponding range is $c_{LG} \in$

$[5.7 \times 10^{18},\ 2.2 \times 10^{19}]$.

What Problem 1 now requires. Equation (73) shows that the derivation of M_* from first principles within PST reduces to a single calculable quantity: the gradient coefficient c_{LG}, which appears in the Landau-Ginzburg functional (33) as the coefficient of the configuration-space gradient term. Equation (C) relates c_{LG} to the measured gravitational constant G up to the normalization κ of the compact fibre projection integral. Once κ is evaluated from the kernel K, equation (74) becomes a derivation rather than an inversion, and M_* is a parameter-free prediction of PST expressed entirely in G, m_h, and v. Computing κ is precisely what Open Problem 1 of Section 14 requires.

10.5. Numerical prediction: $d_0 \approx 7$ nm

Substituting $\ell_* = 1.578 \times 10^{-19}$ m and $\ell_P = 1.616 \times 10^{-35}$ m into equation (66):

$$\begin{aligned} \ell_*^5 &= (1.578)^5 \times 10^{-95}\ \mathrm{m}^5 \ \approx\ 9.74 \times 10^{-95}\ \mathrm{m}^5 \\ \ell_P^2 &= (1.616)^2 \times 10^{-70}\ \mathrm{m}^2 \ \approx\ 2.61 \times 10^{-70}\ \mathrm{m}^2 \\ d_0^3 &= \frac{9.74 \times 10^{-95}}{2.61 \times 10^{-70}} \ \approx\ 3.73 \times 10^{-25}\ \mathrm{m}^3 \\ d_0 &= \left(3.73 \times 10^{-25}\right)^{1/3} \ \approx\ 7.2 \times 10^{-9}\ \mathrm{m} \ =\ 7.2\ \mathrm{nm} \end{aligned} \tag{76}$$

The PST prediction is $d_0 \approx 7$ nm for $M_* = 1.25$ TeV. Allowing M_* to range over 1.0–1.5 TeV gives $d_0 \in [6, 10]$ nm.

The coefficient α_G is now fully determined:

$$\alpha_G \ =\ \frac{\ell_P^2}{d_0^2} \ =\ \frac{(1.616 \times 10^{-35})^2}{(7.2 \times 10^{-9})^2} \ \approx\ 5 \times 10^{-54} \tag{77}$$

This is 54 orders of magnitude below unity, confirming that the provisional claim in Section 7 was quantitatively wrong. The value is not, however, arbitrary. It encodes the same ratio that separates the Planck scale from the electroweak scale: $\alpha_G = (\ell_P/d_0)^2 = (\ell_P/\ell_*)^2 \cdot (\ell_*/d_0)^2$. In PST's framework, the compactification of three substrate dimensions at nanometre scale is the geometric origin of the electroweak hierarchy, the structural reason gravity is so much weaker than the electroweak interaction. The weakness of gravity and the smallness of α_G are the same fact, viewed from the projected geometry and from the precausal substrate respectively.

10.6. Casimir detectability at each separation regime

With $d_0 \approx 7$ nm, the Casimir correction of equation (57) becomes a concrete numerical prediction at each accessible plate separation.

At $d = 160$ nm, the baseline of existing precision measurements [14]:

$$\left(\frac{d_0}{d}\right)^2 \Bigg|_{d=160\,\mathrm{nm}} \ =\ \left(\frac{7}{160}\right)^2 \ \approx\ 1.9 \times 10^{-3} \tag{78}$$

The PST correction is $|\xi| \times 0.19\%$. This lies below the current 1% precision floor; the theory is therefore consistent with all existing data. The prediction is not contradicted: the predicted signal is below the current sensitivity, not above it.

At $d = 50$ nm, accessible to torsion-pendulum and AFM-based Casimir experiments currently under development:

$$\left.\left(\frac{d_0}{d}\right)^2\right|_{d=50\,\mathrm{nm}} = \left(\frac{7}{50}\right)^2 \approx 0.020 \tag{79}$$

The correction is $|\xi| \times 2\%$, above the 1% current floor and detectable at next-generation sensitivity ($\lesssim 0.5\%$ systematic uncertainty).

At $d = 20$ nm, near the practical lower limit for parallel-plate measurements:

$$\left.\left(\frac{d_0}{d}\right)^2\right|_{d=20\,\mathrm{nm}} = \left(\frac{7}{20}\right)^2 \approx 0.12 \tag{80}$$

A $|\xi| \times 12\%$ excess over standard QFT is unambiguous: it exceeds finite conductivity corrections (typically 2–5% at these separations), surface roughness corrections ($\lesssim 1\%$ for well-prepared surfaces), and thermal corrections (subdominant below 100 nm at room temperature [15]).

The distinguishing experimental signature throughout is the power law. Standard QFT and all known material corrections modify the coefficient of d^{-4} or introduce d^{-3} thermal terms; none contribute at order d^{-6}. The d^{-6} PST term is therefore separable by fitting measurements at multiple separations, and the d^{-6}/d^{-4} ratio grows as plate separation decreases, providing an unambiguous directional signature. Substituting $d_0 \approx 7$ nm explicitly into equation (57):

$$\delta P_{\mathrm{PST}}(d) = -\frac{\xi\,\pi^2\hbar c \times (7\ \mathrm{nm})^2}{240\,d^6} \tag{81}$$

The coefficient ξ remains to be computed from the modal kernel K; measuring the d^{-6} component at two separations via the ratio test (59) determines d_0 without requiring knowledge of ξ, and a subsequent fit to the absolute magnitude determines ξ. The two measurements together overconstrain the two parameters, providing an internal consistency check.

10.7. Falsifiability and the experimental programme

The PST Casimir prediction satisfies strict experimental falsifiability. The prediction is refuted if:

1. Precision measurements at $d = 50$ nm agree with standard QFT to better than 1%. This constrains $d_0 < 5$ nm, pushing the implied M_* above approximately 2 TeV, inconsistent with the electroweak condensate identification.

2. Multi-separation ratio tests (equation (59)) in the 20–50 nm range show no d^{-6} component at the 0.5% level.

3. A non-d^{-6} power law is detected, inconsistent with PST's even-in-d_0/d expansion (any

power law intermediate between d^{-4} and d^{-6} would falsify the theoretical structure, not merely constrain d_0).

The prediction is confirmed if:

1. A d^{-6} component is detected at separations below 50 nm, with amplitude consistent with $d_0 \in [6, 10]$ nm.
2. The ratio test between separations $d_1 \approx 20$ nm and $d_2 \approx 160$ nm yields $(d_0/d_1)^2 - (d_0/d_2)^2$ consistent with $d_0 \approx 7$ nm.
3. The fitted d_0 is consistent with $\ell_* = \hbar c/M_*$ for an M_* compatible with LHC measurements of the electroweak new-physics threshold.

Current experiments [14, 15] at 160–750 nm agree with standard QFT to $\sim 1\%$, consistent with $d_0 < 16$ nm from equation (58). The PST prediction $d_0 \approx 7$ nm satisfies this bound. The theory will be testable (and potentially confirmed or refuted) by the next generation of sub-50 nm Casimir experiments currently in preparation in several groups worldwide.

The theoretical relationship between d_0 and ℓ_* is not just a claim: it is a self-consistency test. If experiments find $d_0 = 7$ nm, equation (66) implies $M_* = (\ell_P^2/d_0^3)^{1/5} \cdot \hbar c \approx 1.25$ TeV. This is a prediction for the LHC and its successors: PST expects new physics at approximately 1.25 TeV associated with the full structure of the modal condensate, the precausal origin of the electroweak sector. The two predictions (Casimir d_0, collider M_*) are linked by equation (66) and must be mutually consistent; any inconsistency would refute not merely the numerical prediction but the dimensional reduction mechanism at the core of Section 10.

A note on the scope of falsifiability is warranted. Ruling out $d_0 \approx 7$ nm alone does not break the dimensional-reduction mechanism: it relocates M_* to a different energy scale, keeping the d^{-6} correction structure intact. The prediction acquires its full force only when both channels are tested independently: the d^{-6} Casimir exponent confirms the dimensional-reduction origin of the correction (condition 3 above would rule this out entirely regardless of d_0), and the implied M_* from a measured d_0 must be consistent with collider bounds. If future Casimir experiments find no d^{-6} component at any separation, the dimensional-reduction mechanism (not merely the numerical value of d_0) is falsified.

10.8. Layman's Summary

The substrate grain d_0 is not a free parameter adjusted to fit data; this section derives its value from first principles. As the substrate instantiates, higher dimensions are compactified into a length set by the modal coherence of the threshold, and calculation yields $d_0 \approx 7\,\text{nm}$, a quantitative prediction with no experimental input. At 50 nanometre plate separation in a Casimir experiment this produces a correction of roughly 2 per cent; at 20 nanometres the correction reaches 12 per cent, unambiguous against any competing effect. Detecting the correction at the predicted scale would fix the fundamental modal energy at approximately 1.25 TeV, within reach of particle accelerators.

11. Cosmology: Expansion and Acceleration from Ongoing Sublimation

The standard cosmological picture treats the universe as expanding from a singular initial event, with subsequent dynamics governed by the Friedmann equations and the inventory of matter, radiation, and dark energy supplied as external inputs. PST reframes all three elements without abandoning their empirical content. Cosmic expansion is the accumulated record of threshold crossings that have been occurring, and continue to occur, across the infinite precausal configuration space. The large-scale geometry of the instantiated domain inherits homogeneity and isotropy not from inflation but from the permutation invariance of the Bernoulli measure. And cosmic acceleration is forced, not merely permitted, by the thermodynamics of ongoing instantiation, which simultaneously fixes the sign of the cosmological constant identified as a tension in Section 7.

11.1. The Cosmological Principle as a theorem of the Bernoulli measure

The precausal configuration space is equipped with the Bernoulli product measure $\mu = \bigotimes_{a\in D} \mathrm{Bern}(1/2)$, equation (10), whose defining property is permutation invariance: any finite permutation of the index set D leaves μ unchanged. No locus in configuration space is preferred over any other.

Let j_τ denote the rate of modal threshold crossings per unit proper volume at a point of the instantiated domain $\mathcal{G}$. A spatial variation $j_\tau(\mathbf{x}) \neq j_\tau(\mathbf{y})$ for $\mathbf{x} \neq \mathbf{y}$ would require that the pre-image configurations $\Phi^{-1}(\mathbf{x})$ and $\Phi^{-1}(\mathbf{y})$ are assigned different weight by μ. But μ's permutation invariance assigns equal weight to all configurations related by a finite permutation of loci, and the functor Φ maps isomorphic configurations to isomorphic geometry. No positional preference is available at the precausal level; any apparent variation in j_τ would require a distinguished locus in the substrate, which μ forbids. Isotropy of j_τ follows by the same argument applied to directional labels.

The result is a theorem:

Theorem (Cosmological Principle). The instantiation rate j_τ is spatially uniform and isotropic throughout $\mathcal{G}$ at all scales large relative to d_0. The Cosmological Principle, in standard cosmology a postulate, justified empirically by CMB uniformity [63], is a structural consequence of the permutation invariance of μ.

It follows immediately that the large-scale geometry of $\mathcal{G}$ is described by the Friedmann–Robertson–Walker metric [60]:

$$ds^2 = -c^2\,dt^2 + a(t)^2 \left[\frac{dr^2}{1 - kr^2} + r^2\,d\Omega^2\right], \tag{82}$$

as the unique homogeneous and isotropic solution to the Einstein field equations derived in Section 7, with spatial curvature index k determined by the energy content at the instantiation epoch.

11.2. Horizon and flatness problems dissolved

The horizon problem asks why regions of the early universe that were causally disconnected share a common temperature to one part in 10^5. The flatness problem asks why the spatial curvature parameter Ω_k was tuned to within one part in 10^{60} of zero at the Planck epoch. Standard inflation [50, 51] resolves both by an exponential expansion that carries a single causal patch to the scale of the observable universe.

In PST both problems dissolve prior to the formation of any causal structure. Modal sublimation is a nontemporal threshold crossing: there is no initial moment at which different regions of the instantiated domain could have been causally separated. Homogeneity and isotropy are not consequences of any physical process propagating through $\mathcal{G}$; they are structural properties of the precausal configuration space imprinted on $\mathcal{G}$ at sublimation by the permutation invariance of μ. The uniformity of the CMB [66, 63] is not evidence for inflation; it is direct observational evidence for the permutation invariance of the Bernoulli measure.

The flatness problem also dissolves. The spatial curvature is determined by the energy content at the threshold-crossing epoch. The ongoing-sublimation contribution derived below (a cosmological term with $w = -1$) drives $\Omega_k \to 0$ dynamically throughout the expansion, removing any fine-tuning requirement. No inflaton field or separate scalar sector is needed.

11.3. Ongoing instantiation: why the configuration space does not dilute

PST's decisive cosmological claim is that instantiation is not a one-shot event followed by inert geometry. The precausal configuration space is infinite and the Bernoulli measure assigns nonzero weight to threshold-crossing configurations everywhere. Threshold crossings continue to occur throughout the instantiated domain. The instantiation rate j_τ is constant by the Cosmological Principle theorem above; the energy released per crossing event is $\Delta\mathcal{F} = \mathcal{F}_{\text{pre}} - \mathcal{F}\big|_{\mathcal{V}} > 0$, the descent from the high-tension pre-threshold configuration to the vacuum manifold $\mathcal{V}$ (the sombrero valley of Section 3). The effective energy density of ongoing instantiation is therefore

$$\rho_{\text{inst}} = \frac{j_\tau \, \Delta\mathcal{F}}{c^2}. \tag{83}$$

The key question is whether ρ_{inst} dilutes as the universe expands.

The answer is that it does not. ρ_{inst} depends on two factors: the rate j_τ and the energy quantum $\Delta\mathcal{F}$. By the Cosmological Principle theorem, j_τ is fixed by the permutation invariance of μ and carries no dependence on the size or expansion rate of the instantiated domain. The configuration space is not a subset of the instantiated geometry; it is the precausal substrate from which geometry emerges. Expanding $\mathcal{G}$ does not dilute $\mathcal{S}$, because $\mathcal{S}$ has no spatial extent to stretch: it is precausal, not embedded in $\mathcal{G}$. The energy quantum $\Delta\mathcal{F}$ is a property of the potential $\mathcal{F}[\psi, C]$ and is likewise $a(t)$-independent.

Therefore:

$$\dot{\rho}_{\text{inst}} = 0. \tag{84}$$

11.4. Equation of state and the sign of the cosmological constant

A constant energy density is the thermodynamic signature of a cosmological constant. The covariant energy-conservation equation,

$$\dot{\rho} + 3H(\rho + P) = 0, \tag{85}$$

with $\dot{\rho}_{\text{inst}} = 0$ and $H \neq 0$, forces the pressure to satisfy:

$$P_{\text{inst}} = -\rho_{\text{inst}}, \qquad w \equiv \frac{P_{\text{inst}}}{\rho_{\text{inst}}} = -1. \tag{86}$$

The sign of ρ_{inst} is positive: threshold crossings release tension energy (the configuration descends from the high-tension pre-threshold state to the sombrero valley), so $\Delta\mathcal{F} > 0$ and thus $\rho_{\text{inst}} > 0$. The effective cosmological term contributed by ongoing sublimation is

$$\Lambda_{\text{PST}} = 8\pi G\, \rho_{\text{inst}} > 0. \tag{87}$$

This is positive, de Sitter, not anti-de Sitter. The tension identified in Section 7, equation (32), is thereby resolved: the sombrero valley $\mathcal{F}\big|_{\mathcal{V}} < 0$ would naively give an AdS vacuum from the static vacuum energy alone, but the ongoing-sublimation term $\rho_{\text{inst}} > 0$ is the dominant positive contribution. The sign of Λ is not a free parameter; it is forced by the direction of every threshold crossing (downhill in $\mathcal{F}$, with $\Delta\mathcal{F} > 0$ by definition) and by the impossibility of diluting the configuration-space reservoir.

11.5. Friedmann equations and cosmic acceleration

Substituting the PST stress-energy inventory, pressureless matter, radiation, and the ongoing-sublimation term (ρ_{inst}, $P_{\text{inst}} = -\rho_{\text{inst}}$), into the Einstein field equations of Section 7 with the FRW metric (82) yields the standard Friedmann equations:

$$H^2 = \frac{8\pi G}{3}\Big(\rho_m + \rho_r + \rho_{\text{inst}}\Big) - \frac{kc^2}{a^2}, \tag{88}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\Big(\rho_m + 2\rho_r - 2\rho_{\text{inst}}\Big), \tag{89}$$

where the second equation uses $\rho + 3P = (\rho_m + \rho_r + \rho_{\text{inst}}) + 3(P_m + P_r + P_{\text{inst}}) = \rho_m + 2\rho_r - 2\rho_{\text{inst}}$, with $P_m = 0$, $P_r = \rho_r/3$, and $P_{\text{inst}} = -\rho_{\text{inst}}$. In the instantiation-dominated epoch, $\rho_m, \rho_r \ll \rho_{\text{inst}}$, and equation (89) reduces to

$$\frac{\ddot{a}}{a} \approx \frac{8\pi G}{3}\rho_{\text{inst}} > 0, \tag{90}$$

v9.9

giving late-time acceleration with no additional assumption. The role of ρ_{inst} in PST is structurally identical to Λ in ΛCDM, but the interpretation differs: in ΛCDM, Λ is a free parameter inserted by hand; here, $\Lambda_{\text{PST}} = 8\pi G\,\rho_{\text{inst}}$ is determined, in principle, by the rate j_τ and the tension quantum $\Delta\mathcal{F}$, both of which are properties of the precausal structure. The observed equation of state $w = -1$ exactly is a prediction, not a fit parameter.

11.6. The rate problem

Computing j_τ from first principles requires integrating the threshold-crossing rate over the Bernoulli measure:

$$j_\tau = \int_{\mathcal{S}} \mathbf{1}\big[T(C) \geq \tau\big]\, d\mu(C), \tag{91}$$

where $\mathbf{1}[\cdot]$ is the indicator for configurations at or above the modal threshold. This integral is not presently calculable: the effective restriction of μ to near-threshold configurations, where $\Delta\mathcal{F}$ is well-defined, requires a renormalisation prescription analogous to the Casimir energy calculation of Section 9. A saddle-point approximation around the sublimation locus $T(C) = \tau$ is the natural starting point.

This is listed as an open problem in Section 14 (Problem 2*b*). The observed value $\Lambda_{\text{obs}} \approx 1.1 \times 10^{-52}\ \text{m}^{-2}$ provides a target: $j_\tau \cdot \Delta\mathcal{F} = \Lambda_{\text{obs}}\, c^4/8\pi G \approx 5.4 \times 10^{-10}\ \text{J}\,\text{m}^{-3}$. A successful computation of j_τ would determine Λ_{PST} from first principles and, together with the Casimir prediction for d_0, constitute a second independent test of the dimensional-reduction mechanism at the core of Section 10.

11.7. Observational commitments

Beyond the standard ΛCDM predictions, PST makes the following specific commitments:

1. **CMB uniformity.** Permutation invariance of μ guarantees homogeneity and isotropy prior to any causal structure. Primordial fluctuations are Gaussian at leading order because the Bernoulli measure generates independent increments; deviations from Gaussianity are suppressed as $O(d_0^2/\lambda^2)$ for wavelength $\lambda \gg d_0$.

2. **Big Bang nucleosynthesis.** PST does not alter the microphysics of BBN. The weak-interaction rates, neutron-to-proton ratio, and light-element abundances depend on the QM sector, which emerges as a projection of the precausal structure (Section 5) and is indistinguishable from standard QM at energies well below $M_* \approx G/d_0^2$.

3. **Equation of state.** The instantiation-dominated epoch gives $w = -1$ exactly, consistent with current constraints [63] and distinguishable from dynamical dark energy ($w \neq -1$) by forthcoming surveys (Euclid, DESI).

4. **Spatial flatness.** The driving of $\Omega_k \to 0$ by ρ_{inst} and the absence of initial fine-tuning predict $\Omega_k = 0$ to high precision, consistent with current CMB bounds [63].

5. **Cosmic acceleration.** Equation (90) gives late-time acceleration consistent with Type Ia supernova data [61, 62] without introducing new fields beyond those required by the

precausal structure.

11.8. Comparison with steady-state cosmology

The continuous creation of energy density in Bondi–Gold–Hoyle steady-state cosmology [64, 65] is the closest historical precedent for PST's ongoing instantiation: both predict $\dot{\rho} = 0$ and $w = -1$ as a consequence of an ongoing creation process. The differences are decisive.

Mechanism. Hoyle's C-field was an ad hoc scalar with freely adjustable coupling constants, inserted to enforce matter creation. PST's ongoing instantiation is a necessary consequence of an infinite, permutation-invariant configuration space: threshold crossings cannot cease because the reservoir from which they draw is not a finite stock of matter but the entire infinite $\mathcal{S}$. Termination would require a privileged locus in $\mathcal{S}$ at which μ vanishes, which permutation invariance forbids.

What is created. Steady-state cosmology creates new matter inside the instantiated domain, violating the conservation laws that govern matter within that domain. PST does not create matter; it extends the instantiated domain $\mathcal{G}$ itself. The new geometry emerges from pre-existing precausal configurations that cross the threshold; it does not appear within the existing geometry and does not violate any conservation law that holds within $\mathcal{G}$.

The CMB. The steady-state model predicted no CMB; its refutation by Penzias and Wilson [66] was decisive. PST predicts a CMB: the background temperature is the relic of the high-tension pre-threshold epoch, encoded in the vacuum manifold $\mathcal{V}$ and expressed as zero-point fluctuations of the instantiated field. The CMB is evidence for PST, not against it.

Structure formation. Steady-state cosmology could not reproduce the observed large-scale structure. PST predicts a near-scale-invariant Gaussian spectrum of primordial perturbations from the Bernoulli measure, consistent with Planck CMB observations [63] and the transfer-function structure of ΛCDM.

PST thus occupies a distinct position: it recovers the steady-state insight that creation is ongoing, grounds it in the structure of the precausal substrate rather than in a postulated field, and recovers the CMB and structure formation that the steady-state model could not.

11.9. Layman's Summary

PST reframes cosmology from the ground up. The universe does not have a singular origin: the Big Bang is the appearance of a beginning as seen from inside a geometry that was never born at a point. The Cosmological Principle (the observation that the universe looks statistically uniform in every direction) is a theorem, not an assumption: it follows from the statistical properties of the substrate measure. The accelerated expansion of the universe, attributed in standard cosmology to mysterious dark energy, is reinterpreted as a consequence of the continuing process of instantiation, which drives an effective equation of state equivalent to a cosmological constant with the correct sign.

12. Relation to Existing Formalisms

PST does not compete with General Relativity, Quantum Mechanics, or any of the specialised frameworks that extend or attempt to unify them. Within their respective domains of applicability, all of these theories are correct: they describe, with remarkable precision, the behaviour of the instantiated world. What PST offers is not a replacement of any of them but a *foundation*: a precausal substrate from which each emerges, and within which the questions each leaves unanswered find principled answers. This section examines each major existing framework in turn, identifying both what it correctly describes and precisely where PST deepens or extends it.

12.1. General Relativity

General Relativity [1, 10] is the most successful classical theory of gravitation ever formulated. It correctly describes the curvature of spacetime by matter and energy, the precession of planetary perihelion, the deflection of light by mass, the existence and properties of black holes, the expansion of the universe, and the propagation of gravitational waves, confirmed directly by LIGO [46]. Its predictive precision over more than a century of increasingly demanding tests places it among the most thoroughly verified theories in all of science.

What General Relativity cannot explain is its own foundation. It assumes a smooth Lorentzian manifold without providing any account of why such a manifold exists or why it has the Lorentzian rather than Riemannian signature. It takes the stress-energy tensor as an external input without deriving it. It introduces the gravitational constant G without fixing its value or explaining its dimensional role. The geodesic equation, which describes how bodies move, specifies which paths are available but provides no explanation for why a body occupies one path rather than another: the choice of geodesic remains an initial condition imposed from outside. The initial singularity, required by the Penrose-Hawking singularity theorems [2, 3], terminates the theory at the beginning of cosmic time and leaves the origin of the universe outside its reach. Diffeomorphism invariance, the symmetry under arbitrary coordinate transformations, is imposed as a postulate rather than derived.

In PST, each of these lacunae is filled. General Relativity emerges as the conservation law of the tension projection $\Pi(T(C))$ across the functor Φ: the field equations $G_{\mu\nu} = 8\pi G \cdot \Pi(T(C))$ express the fact that total asymmetric tension is preserved across the sublimation threshold, distributed simultaneously as curvature and as stress-energy. Diffeomorphism invariance is not postulated but derived: it corresponds to the non-uniqueness of the realisation map ρ, which assigns points of the manifold to properties of the precausal domain without a canonical choice. The constant G is a unit conversion factor introduced by Π, not a fundamental property of the substrate. The geodesic gap is closed by the sombrero topology of the vacuum manifold: orbital motion is not an initial condition but the unique stable state of a configuration on $\mathcal{V} \cong S^1$. The initial singularity is dissolved because loci are not primitive in PST: there is no single point at which reality began.

12.2. Quantum Mechanics and Quantum Field Theory

Quantum Mechanics [23, 24] and its relativistic extension, Quantum Field Theory [25], together describe all known non-gravitational physics to extraordinary precision. The Hilbert space formalism, the Born rule, the superposition principle, the uncertainty principle, and the operator algebra of observables all emerge from a remarkably compact mathematical framework. The predictions of QFT, including the anomalous magnetic moment of the electron and the Lamb shift, match experimental measurements to better than one part in 10^{12}.

Yet Quantum Mechanics stands on several postulates whose justification is never given from within the theory. The Hilbert space structure is assumed. The complex-valued probability amplitudes are assumed. The Born rule, connecting amplitudes to measurement probabilities, is assumed. The measurement postulate, specifying that the state vector collapses to an eigenstate upon observation, has no derivation from the Schrödinger equation and has been the source of deep interpretational controversy for a century [24]. QFT further requires a fixed background spacetime on which its fields are defined; it has no consistent formulation in strongly curved or dynamical geometries. The vacuum energy density predicted by QFT from zero-point fluctuations of all fields exceeds the observed cosmological constant by approximately 120 orders of magnitude, the worst quantitative discrepancy in the history of theoretical physics.

In PST, Quantum Mechanics and Quantum Field Theory are derived at two distinct levels of the projected structure. QM is the near-threshold physics: the regime $\varepsilon(C) \approx 0$ in which the order parameter has not settled firmly into either minimum. The wave function corresponds to the distribution of ψ across the nearly flat landscape near τ. Superposition is not a metaphysical puzzle but a structural feature of configurations genuinely intermediate between the two wells: the order parameter explores both. The measurement problem dissolves: collapse is the settling of ψ from the near-threshold region into one minimum, induced by sufficient coupling to a well-settled macroscopic configuration.

QFT is the further projection: the spectrum of fluctuations of the order parameter about the settled vacuum $\mathcal{V}$, promoted to a Lorentzian field theory by the functor Φ as derived in Section 7. The radial mode is massive with $m_h^2 = 4\varepsilon$; the phase mode is massless; their dynamics are governed by the instantiated action S, not the free energy $\mathcal{F}$. QFT presupposes the Lorentzian action and therefore presupposes the threshold crossing; it is a projected description of the instantiated domain, not a competitor to PST at the precausal level. The open problems in the QFT sector are stated explicitly in Section 7: the Born rule, the spin–statistics connection, and the full interacting vacuum are not yet derived.

The vacuum energy problem is *relocated*: the physical vacuum is not $\psi = 0$ but the sombrero ring $\mathcal{V}$ (Section 7), which relocates where the calculation must be performed. The finite value that emerges from the correct evaluation point is not yet computed from first principles; this is an open problem, related to the rate integral of Section 11.

12.3. Thermodynamics and Black Hole Entropy

The thermodynamics of black holes, initiated by Bekenstein [47] and Hawking [48], established that black holes carry an entropy proportional to their horizon area, $S = A/(4\ell_P^2)$, and a temperature proportional to their surface gravity. These relations hint at a deep connection between geometry, thermodynamics, and information that is not explained by either General Relativity or Quantum Mechanics taken alone. Jacobson [18] later showed that the Einstein field equations could be derived by applying the first law of thermodynamics to a local Rindler horizon, suggesting that gravity is an entropic phenomenon. Verlinde [17] extended this to propose that gravity is an emergent entropic force on the holographic screen of a region.

PST provides a natural interpretation of all these results and locates their origin in the precausal substrate. The entropy of a black hole, in PST terms, is the number of distinct precausal configurations C that project to the same black hole geometry via Φ: it counts the modal degeneracy of the instantiated state. Since the realization map ρ is non-injective, multiple configurations can project to manifolds with the same macroscopic geometry, and the Bekenstein entropy counts precisely this degeneracy. The area law $S \propto A$ follows from the fact that ρ maps internal structure to boundary data: the information content of the projection is encoded on the surface where the metric transitions from interior to exterior. Jacobson's derivation of the field equations is, from the PST perspective, a thermodynamic reformulation of the conservation law $G_{\mu\nu} = 8\pi G \cdot \Pi(T(C))$: both express the same content, one in geometric language and the other in entropic language.

12.4. Causal Set Theory

Causal set theory [4, 5] proposes that at the fundamental level spacetime is discrete: the continuum manifold is an approximation to a locally finite partial order, the causal set, in which each element represents an elementary spacetime event and the partial order represents causal precedence. The continuum metric is to be recovered in an appropriate limit. This approach takes seriously the need for a background-independent, Lorentz-covariant discretisation of spacetime and has produced significant results on the emergence of the cosmological constant from causal set dynamics.

PST shares with causal set theory the rejection of the continuum manifold as a fundamental given, but the two frameworks differ on what is primitive. In causal set theory, the causal partial order $\leq$ is the primitive structure from which the continuum metric is to be recovered. In PST, both the metric and the causal order are derived from the precausal substrate. The causal order, as established in Section 5, has no pre-image in the precausal category $\mathbf{S}$: it is generated anew by the Lorentzian signature of the metric produced by Φ. This means that causal set theory, in taking the partial order as its starting point, is already working within the instantiated domain. Its primitive is a structure that PST derives. The discrete causal set arises naturally near the modal threshold, where configurations are only marginally instantiated and the realisation map ρ is sensitive to individual pairwise tensions: this is the precausal origin of causal discreteness, and it appears only in the near-threshold regime. At large $\varepsilon(C)$, the continuum limit is recovered, consistent with causal set theory's own continuum approximation conjecture.

12.5. Loop Quantum Gravity

Loop Quantum Gravity [6, 7] quantizes geometry directly, without first identifying a background metric. Working with the Ashtekar connection variables, it produces a discrete spectrum of geometric operators: area and volume have quantized eigenvalues at the Planck scale, and the states of the gravitational field are described by spin networks, graphs whose edges carry SU(2) representations encoding quantized areas. LQG has made substantial progress on the problem of quantum geometry and has proposed a resolution of the initial singularity through a bounce in loop quantum cosmology.

Several foundational questions remain open within LQG. The Barbero-Immirzi parameter γ, which appears in the definition of the Ashtekar-Barbero connection, is not fixed by the theory and must be set by matching the black hole entropy calculation. The relation between spin networks and classical geometry is understood at the level of expectation values of geometric operators but not at the level of the semiclassical limit. The problem of time, the difficulty of interpreting a Hamiltonian constraint that generates no evolution, remains unresolved.

PST relates to LQG at the level of the near-threshold regime. Spin networks are the appropriate description of a highly non-uniform, near-threshold tension configuration: the graph structure encodes the topology of connections between marginally-instantiated sub-configurations, and the SU(2) labels encode the quantized angular momentum $L = n\hbar$ arising from the U(1) vacuum symmetry. The Barbero-Immirzi parameter is, in PST terms, related to τ_{null}: the ratio of the null threshold to the full modal threshold sets the boundary between timelike and spacelike separations and thus fixes the overall scale of the quantum geometric operators. The problem of time dissolves because time is not a primitive in PST: it is coinstantiated with geometry through Φ, so the Hamiltonian constraint is not a constraint on evolution but on the consistency of the projection.

12.6. String Theory and M-Theory

String Theory [26, 27] is the most elaborately developed framework for unifying quantum mechanics with gravity. By replacing point particles with one-dimensional strings propagating in a target spacetime, it naturally produces a massless spin-2 excitation identified with the graviton, finite scattering amplitudes at all orders, and, through the superstring spectrum, the gauge bosons and fermions of a potential unified theory. M-Theory, the conjectured eleven-dimensional theory that unifies the five consistent superstring theories, extends this further. The framework has produced deep mathematical structures, including mirror symmetry, D-branes, and the AdS/CFT correspondence.

The foundational difficulties of string theory are equally significant. The theory requires ten spacetime dimensions for mathematical consistency (eleven for M-theory), and the six extra dimensions must be compactified on a Calabi-Yau manifold to recover four macroscopic spacetime dimensions. The choice of Calabi-Yau manifold is not determined by the theory: the string landscape contains an estimated 10^{500} distinct vacua, each corresponding to a different compactification and a different set of low-energy physics. The anthropic principle is invoked to explain why we inhabit one vacuum rather than another. More fundamentally, string theory still requires a background spacetime through which strings propagate; it is not truly background-independent. The question of why strings, rather than some other fundamental object, is unanswered.

PST and string theory agree that physical reality is higher-dimensional than it naively appears, but they differ on every other structural point. In string theory, extra dimensions are a mathematical consistency requirement, imposed from the outside and then hidden by compactification; in PST, higher dimensionality is the natural state of an instantiated geometry close to the modal threshold, and the reduction toward four dimensions is driven by the growth and settling of the tension distribution, not by geometric compactification. PST has no landscape: the tension configuration C that produces our universe is a specific point in the modal potential, not a choice from an exponentially large set of equally valid alternatives. The question of why strings is not asked in PST because strings are not primitive: the extended objects of high-energy physics are near-threshold tension structures whose one-dimensional character reflects the topology of the configuration space at marginal instantiation.

12.7. Kaluza-Klein Theory and Extra Dimensions

Kaluza [28] and Klein [29] were the first to propose that extra spatial dimensions could unify fundamental forces. Adding a fifth dimension to General Relativity, and compactifying it to a circle of Planck-scale radius, produces four-dimensional gravity plus electromagnetism from a single five-dimensional geometric framework. The U(1) gauge symmetry of electromagnetism arises as the isometry group of the compactified circle. This is one of the most elegant results in the history of theoretical physics: a symmetry of fundamental physics derived from the geometry of a hidden dimension.

The relationship between Kaluza-Klein and PST is particularly close. The U(1) isometry of Kaluza-Klein's compactified circle corresponds precisely to the U(1) symmetry of the vacuum manifold $\mathcal{V} \cong S^1$ in PST. In PST, this symmetry is not a geometric property of a hidden extra dimension but a structural property of the modal potential functional in the instantiated regime. The circle S^1 that Kaluza-Klein compactifies as an extra spatial dimension is, from the PST perspective, the vacuum manifold of the sombrero potential: not an additional geometric direction alongside the four macroscopic ones but the topological structure of the degenerate vacuum that gives rise to the conserved U(1) charge identified as electromagnetism. Kaluza-Klein correctly identifies the U(1) symmetry and its geometric interpretation; PST identifies its precausal origin. In PST, the extra dimension does not need to be compactified to a microscopic scale by hand because it was never an independent spatial direction; it is the phase direction of the complex order parameter, which is intrinsically a modal degree of freedom rather than a spatial one.

12.8. Large Extra Dimensions and Braneworld Models

The Arkani-Hamed-Dimopoulos-Dvali (ADD) model [30] proposes that the apparent weakness of gravity relative to other forces, the hierarchy problem, could be explained if gravity propagates in extra dimensions of submillimeter size while the Standard Model fields are confined to a four-dimensional hypersurface. The extra dimensions lower the fundamental scale of gravity to the TeV range, making quantum gravity effects accessible at collider energies. Randall-Sundrum models [31] go further, proposing a warped extra dimension between two branes, from which they derive the hierarchy without requiring large extra dimensions.

Both the ADD model and the Randall-Sundrum framework share with PST the prediction that extra-dimensional effects are observable at short distances and high energies, but both require background structures that PST derives. The ADD model posits a fixed flat background with extra dimensions of fixed size; the Randall-Sundrum model posits a specific five-dimensional anti-de Sitter geometry as the background. In both cases the number of dimensions, the size of the extra dimensions, and the geometry of the bulk are inputs, not outputs. PST provides a principled answer to the question that both models leave open: why are there extra dimensions at all, and why do they have the structure they do? In PST, the extra dimensions of the early instantiation are not background features but the natural consequence of a tension configuration close to the modal threshold, where the order parameter has not yet settled and the geometry is still high-dimensional. Their subsequent reduction to four macroscopic dimensions is driven by the growth of instantiation, not by stabilisation mechanisms or brane tensions imposed from outside.

12.9. Causal Dynamical Triangulations

Causal Dynamical Triangulations (CDT) [32] is a non-perturbative approach to quantum gravity that computes a sum over causal geometries, discretised as simplicial complexes with a built-in causal structure. One of its most striking results is the observation that the spectral dimension of the resulting quantum geometry is scale-dependent: it approaches two at the Planck scale and recovers four at large scales. This spontaneous dimensional reduction is obtained from first principles, without assuming four dimensions as a background, and it represents one of the clearest quantitative predictions of any approach to quantum gravity.

PST predicts a qualitatively identical pattern: the dimensionality of instantiated geometry is high near the modal threshold and reduces toward four as the geometry grows and the tension distribution settles. The PST mechanism is explicit and derivable: the order parameter $r_0 = \sqrt{\varepsilon/(2b)}$ grows with the excess tension $\varepsilon(C) = T(C) - \tau$, and as r_0 grows the geometry explores progressively fewer independent directions in configuration space, with the U(1) vacuum manifold providing the final settled structure at four dimensions. CDT's dimensional flow and PST's dimensional reduction thus agree on the direction and the endpoint; they differ on the mechanism. CDT's result is a quantum averaging effect over discrete triangulations, with causality imposed as a constraint at the level of the sum over histories. PST's result follows from the modal potential, without requiring a sum over geometries or an imposed causal constraint, since causality itself is derived. PST provides CDT with the precausal origin of the causal constraint that CDT must impose.

12.10. The Holographic Principle and AdS/CFT

The holographic principle, proposed by 't Hooft [33] and developed by Susskind, Bousso, and others, states that the maximum information content of a gravitational region is not proportional to its volume but to the area of its boundary. The Bekenstein-Hawking entropy $S = A/(4\ell_P^2)$ is the quantitative expression of this bound. The AdS/CFT correspondence of Maldacena [34] realises the holographic principle explicitly: the string theory in an anti-de Sitter bulk is exactly equivalent to a conformal field theory on its boundary, providing a concrete duality between a D-dimensional gravitational theory and a $(D-1)$-dimensional non-gravitational one.

PST offers a precausal foundation for holography. The reduction in effective dimensionality as instantiated geometry grows is the precausal process that gives rise to holographic encoding. The realization map $\rho : D \to M$ assigns to each property in the precausal domain a point in the instantiated manifold. At early instantiation, many independent directions in configuration space project into the manifold; as geometry grows and the order parameter settles, progressively fewer independent directions remain, and the information content of the bulk is compressed onto the boundary of the settled region. The holographic bound $S \leq A/(4\ell_P^2)$ is thus not a fundamental principle but a derived property of the projection operator Π: it reflects the fact that ρ maps the internal structure of the precausal domain onto boundary data, and that the amount of boundary data scales with area rather than volume because the projection collapses interior degrees of freedom. AdS/CFT in PST is a specific realisation of this general mechanism in the particular geometry of an anti-de Sitter bulk, arising from a tension configuration with a specific asymptotic structure.

12.11. Non-Commutative Geometry

Non-commutative geometry, developed by Connes [49], replaces the classical notion of a smooth manifold with a spectral triple (A, H, D), consisting of an algebra A of functions, a Hilbert space H on which they act, and a Dirac operator D that encodes the metric geometry. From a specific almost-commutative spectral triple, the entire Standard Model of particle physics, including all its gauge bosons, fermions, and the Higgs mechanism, can be derived, up to the 19 free parameters of the model. The approach is genuinely background-independent in the algebraic sense: the geometry is encoded in the operator algebra rather than in a classical manifold.

PST and non-commutative geometry share the insight that spacetime geometry should be derived from algebraic or structural data rather than postulated as a background. The key difference is the level of foundation. Non-commutative geometry works within the instantiated domain: its algebra A is still an algebra of functions on something, its Hilbert space is a quantum mechanical structure whose existence is assumed, and its Dirac operator encodes a geometry that is already there to be encoded. PST works from the precausal level: the algebra of distinctions δ and the tension functional T are more primitive than any operator algebra, and the Hilbert space structure of quantum mechanics, including the specific algebra of observables, is to be derived from the projection of the modal potential near threshold. Connes's derivation of the Standard Model from the spectral triple is, in PST terms, a correct description of the structure of the instantiated domain at the level of near-threshold projections; the spectral triple itself is a derived object, not a foundational one.

12.12. Inflationary Cosmology and the Multiverse

Inflationary cosmology [50, 51] proposes that the very early universe underwent a brief period of exponential expansion driven by the potential energy of a scalar field, the inflaton. This resolves several otherwise puzzling features of the standard cosmological model: the near-perfect uniformity of the cosmic microwave background (the horizon problem), the near-exact spatial flatness of the universe (the flatness problem), and the absence of magnetic monopoles. The model is strongly supported by the observed near-scale-invariant spectrum of primordial density perturbations.

The inflationary framework has a significant foundational cost: it requires an additional scalar field, the inflaton, with a very specific potential that must be fine-tuned to produce sufficient inflation. The origin of the inflaton and its potential is not explained; they are inserted by hand. Moreover, in most inflationary models, the field drives inflation eternally in some regions while ending in others, giving rise to an eternal inflation scenario in which our observable universe is one of infinitely many causally disconnected pocket universes, each with potentially different low-energy physics. The resulting multiverse makes the precise values of physical constants in our universe a matter of anthropic selection rather than physical necessity.

PST requires no separate inflaton field. The near-threshold phase of instantiation, in which the order parameter ψ is small and the geometry is high-dimensional, corresponds directly to the inflationary epoch: the rapid growth in the apparent scale of the four-dimensional geometry is not driven by a separate field but by the settling of the tension distribution as geometry grows. The horizon and flatness problems are resolved because the precausal substrate instantiates geometry continuously and everywhere at once; there is no finite initial region that needed to be inflated into the visible universe. The scale-invariant spectrum of density perturbations follows from the near-threshold fluctuation structure of $\delta g_{\mu\nu}$ in equation (51): near τ, the power spectrum of metric fluctuations is nearly scale-invariant because the modal potential is nearly flat across all scales. The multiverse does not arise in PST because loci are not primitive: the substrate instantiates geometry in a single continuous field of instantiation, not in discrete pocket universes. There is no ensemble of separate universes from which ours is selected anthropically; there is one precausal substrate whose single continuous instantiation includes everything that is.

The quantum/gravity divide dissolves in this picture because both theories describe the same structure, the tension functional $T(C)$ projected through Π, from different regimes of the excess tension $\varepsilon(C)$: classical geometry at $\varepsilon \gg 0$, quantum behaviour at $\varepsilon \approx 0$. The apparent incompatibility between General Relativity and Quantum Mechanics is not a conflict between two fundamental theories but a parallax error arising from describing the same underlying structure from different vantage points.

12.13. Spencer-Brown and Distinction-Based Ontologies

G. Spencer-Brown's *Laws of Form* [53] is the closest formal precursor to PST's foundational move. Spencer-Brown begins from a single primitive operation, the drawing of a distinction, and develops from it a calculus of indications that regenerates Boolean algebra and, through re-entry of the form into itself, the recursive structures of self-reference. The act of distinction is taken as logically prior to every other structure: prior to objects, sets, truth values, and identity. Nothing is more primitive; the distinction cannot be grounded in anything that does not already presuppose it.

PST shares this identification of the distinction as primitive but differs in three respects. First, PST's (D, δ) axiomatises the *result* of distinction, the structure of already-differentiated properties, rather than the performative act of drawing one; PST has no observer and no injunction. Second, Spencer-Brown's calculus remains within the instantiated logical domain: the tokens on each side of the distinction are already placed and available to be labelled. PST's property differentiation is pre-logical: the properties have no location, no identity conditions other than their mutual non-identity, and no prior domain in which they are situated. Third, PST extends the formal-logical move into the physical by adding asymmetric tension and the modal threshold; Spencer-Brown's laws contain no mechanism for generating geometry from distinction alone. The genealogical connection is nonetheless real: *Laws of Form* is the prior work that most clearly establishes distinction as an irreducible primitive, and PST is its physical continuation.

Hegel's *Science of Logic* [59] provides an earlier philosophical foreshadowing. The opening dialectic moves from pure being, which has no determinate content, to nothing, and then resolves both in becoming, the first concrete category. The precausal substrate has no geometric or causal content (the analogue of pure being collapsing toward nothing), while modal sublimation is the transition by which indeterminate tension becomes determinate geometry (the analogue of becoming as resolution). PST supplies what Hegel does not: a mathematical mechanism specifying the threshold condition $T(C) > \tau$ and the precise variational form of the transition.

12.14. Wheeler's "It from Bit" and the Mathematical Universe

John Wheeler's "it from bit" thesis [54] is the most influential information-theoretic predecessor to PST. Wheeler's central claim is that every particle, field, and spacetime event derives from binary distinctions, from answers to yes-or-no questions registered by physical apparatus. The universe is participatory: it comes into being through the act of observation.

PST agrees with Wheeler that the fundamental ontology is structural rather than material, and that distinctions are more primitive than things. It departs from Wheeler's participatory account at a foundational point. In PST, the substrate does not require an observer to instantiate; it requires only that internal tension exceed the modal threshold τ. Observation is a derived phenomenon occurring within the instantiated geometry that sublimation produces. Wheeler's "bits" correspond in PST to the complement-asymmetries $T(C) \neq T(\bar{C})$: not binary answers to questions posed by observers but structural facts about the precausal configuration that obtain prior to any apparatus.

Tegmark's Mathematical Universe Hypothesis [55] carries mathematical realism further, proposing that all self-consistent mathematical structures are physically real and that our universe is one such structure selected by its initial conditions. PST converges with Tegmark in grounding physical reality in mathematical structure rather than matter. It diverges on the status of the primitive. For Tegmark, mathematical structure is itself the primitive, and the question of which structure is instantiated is answered anthropically. For PST, mathematical structure is derived from the more primitive (D, δ): the configuration space, the tension functional, and the modal potential are mathematical objects generated by the axioms, not given as brute ontological facts. The question of which configuration is instantiated is not anthropic in PST but structural: the configuration that exceeds τ cannot remain uninstantiated by logical necessity rather than observer selection.

12.15. Structural Realism and Bohm's Implicate Order

Ontic structural realism (OSR), as developed by Ladyman [56] and French [57], holds that what is physically real is not the intrinsic properties of objects but the relational structures that physical theories describe. In the strongest version, there are no objects with intrinsic natures underlying the relations; the relational structure is all there is. This position is motivated by the success of structural descriptions in physics, symmetry groups, state spaces, transformation laws, and by the instability of object-based ontologies under theory change.

PST is a natural ally of OSR, working at the deepest structural level available. Where OSR typically takes the relational structures of existing physical theories as its ontological ground, PST asks what those structures themselves emerge from. The answer is the single relation δ and the complement-asymmetry condition $T(C) \neq T(\bar{C})$. PST is thus OSR carried one step further back: not the structures of physics as the ground but the proto-structural primitive from which those structures are generated. The relational contents of quantum mechanics, general relativity, and the Standard Model are, on this account, projections of a single more primitive relational fact through the operator Π.

Bohm's implicate order [58] provides a complementary comparison. Bohm proposed that the apparent discreteness and separateness of physical objects constitute an explicate order projected from a deeper implicate order in which everything is enfolded into everything else. The holistic character of the implicate order and its projection into local explicate structure closely parallels the PST functor $\Phi : \mathbf{S} \to \mathbf{G}$, which projects the holistic precausal substrate into the locally structured instantiated geometry. The key difference is formalisation: Bohm's implicate order is developed at the level of physical intuition and interpretive framework without a mathematical axiomatisation, and no mechanism specifies when or how the projection occurs. PST provides the structure that Bohm's picture envisions: the precausal substrate is the implicate order, and Φ is the explicate projection, with the bifurcation condition $T(C) > \tau$ specifying precisely when projection occurs.

12.16. Layman's Summary

This section places PST in the landscape of existing physical and philosophical theories. General Relativity is recovered in full: the field equations, the equivalence principle, and the Lorentzian signature are all theorems of PST rather than assumptions. Quantum Mechanics emerges as the description of configurations near the modal threshold, where geometry is not yet fully settled. String theory, loop quantum gravity, causal set theory, and non-commutative geometry are each examined; PST is shown to operate at a deeper level, deriving the background structures that each of these theories takes as its starting point.

13. Conclusion

Precausal Substrate Theory begins from a single observation and follows its consequences without concession. The observation is this: any account of physical reality that takes spacetime, causality, or matter as its starting point has already assumed the very thing it needs to explain. Every existing framework in theoretical physics, however successful, begins in the middle of the story. General Relativity begins with a smooth Lorentzian manifold and derives its dynamics. Quantum Mechanics begins with a Hilbert space and derives its predictions. String Theory begins with strings in a background target space and derives its spectrum. None of these frameworks asks why the starting point exists, why it has the structure it does, or whether that structure is necessary or contingent. PST asks all three questions, and provides answers.

13.1. The central claim, stated precisely

PST proposes that physical reality does not originate within spacetime. It arises from a precausal substrate: a nonspatiotemporal, non-causal domain whose only primitive is the capacity for distinctions to exist, what the theory calls property differentiation. From this single primitive, the entire landscape of physical structure, spacetime, causality, matter, energy, gravity, quantum behaviour, angular momentum, and orbital mechanics, is derived by logical entailment. No step in the derivation introduces a new primitive; each concept emerges necessarily from the preceding one. The chain terminates not arbitrarily but at the topology of the stable vacuum, the point beyond which further structure depends on the specific content of the precausal configuration rather than on its logical form.

The claim is a strong one, and it is important to state what it does and does not assert. PST claims to identify the necessary preconditions for any instantiated reality: the minimal ontological structure from which a Lorentzian manifold, together with its causal order, its matter content, and its gravitational dynamics, follows necessarily. It does not claim to derive the specific numerical values of all physical constants, the specific gauge group of the Standard Model in complete detail, or the specific matter content of our universe. These depend on the particular tension configuration C that produced our universe and require knowledge of C beyond its structural form. What PST derives is the framework within which all such specific facts are intelligible.

13.2. What the emergence chain establishes

The core argument of the paper proceeds in eight steps, each logically entailed by the previous. Property differentiation (D, δ) is the single primitive: the capacity for differences to exist, irreducible because any attempt to explain it presupposes it. From property differentiation, asymmetric tension follows: any configuration of two or more distinct properties carries a structural imbalance, a complement-asymmetry $(T(C) \neq T(\bar{C}))$ that cannot be neutralised within the precausal domain. From asymmetric tension, the modal threshold follows: the tension functional $T(C)$ admits a critical value τ at which the stable uninstantiated position $\psi = 0$ becomes an unstable critical point of the modal potential functional $\mathcal{F}(\psi, C)$. Once $T(C) > \tau$, remaining uninstantiated is logically incoherent, not merely unstable. The resolution of that incoherence is modal sublimation: the direct, nontemporal, nonmechanistic transition of the precausal configuration into instantiated geometry, represented mathematically as the functor $\Phi : \mathbf{S} \to \mathbf{G}$.

The functor Φ produces a Lorentzian manifold (M, g) whose signature $(-, +, +, +)$ is not postulated but derived from the bifurcation of pairwise tensions around the null threshold τ_{null}. Causality, the partial order $p \leq q$ on M, is coinstantiated with the manifold: it has no pre-image in the precausal category $\mathbf{S}$ and arises anew from the metric structure of g. The field equations of General Relativity emerge as the conservation law of the tension projection $\Pi(T(C))$ across the threshold: both curvature and stress-energy are outputs of the same projection, not independent inputs. Angular momentum emerges as the Noether charge of the U(1) symmetry of the vacuum manifold $\mathcal{V} \cong S^1$, deriving the universality of orbital motion from the topology of the degenerate vacuum without any initial condition.

Each step of this chain is mathematically explicit. The modal threshold is defined by the infimum condition (16). The tension-to-metric map is equation (18). The stress-energy tensor is equation (24). The conservation law is equation (25). The Noether charge is equation (36). The chain is not a qualitative narrative but a sequence of mathematical constructions, each of which is both necessary and sufficient for the next.

13.3. The open questions PST closes

Physics and philosophy have accumulated a set of questions that existing frameworks consistently leave unanswered. PST closes each of them, not by providing ad hoc responses but by exhibiting the structural reason why they arise and dissolving them at the level of their formation.

Why is there something rather than nothing? Leibniz's question [42] receives not an arbitrary answer but a principled one. The logical possibility of property differentiation is irreducible: any context in which the question is asked is already a context in which distinctions exist, presupposing differentiation. A substrate bearing irreducible property differentiation cannot remain uninstantiated once the threshold condition is satisfied, because remaining uninstantiated is logically incoherent above τ. Existence is not a brute contingent fact; it is the necessary expression of a logically unavoidable modal imbalance.

What caused the universe to come into being? This question is not unanswerable. It is a category error: it applies the causal relation to the event of universal instantiation, but the causal relation is coinstantiated with the universe and does not exist prior to it. Causality is a derived structure, causality $= f(g) = f(\Phi(C))$; it cannot serve as the explanation for the event that constitutes its own first instantiation. The question does not lack an answer; it lacks a well-formed domain.

Why do General Relativity and Quantum Mechanics appear incompatible? Both theories are correct within their domains. Their apparent incompatibility is a parallax error: they describe the same precausal structure, $T(C)$ projected through Π, from different vantage points. At $\varepsilon(C) \approx 0$, near the modal threshold, the projection is sensitive to small fluctuations in tension, producing indeterminate, wavelike behaviour: this is the quantum regime. At $\varepsilon(C) \gg 0$, far from the threshold, the order parameter is settled and the projection produces a smooth, deterministic geometry: this is the classical relativistic regime. The gap between the two theories is not a gap between two realities but between two limits of a single one.

Why does matter curve spacetime? Because matter and curvature are not two things. Both are projections of the same asymmetric tension $T(C)$ through the operator Π. Einstein's field equations are not a law relating two independent ontological kinds; they are a conservation statement expressing that the total tension content of the precausal configuration is preserved across the sublimation threshold, distributed simultaneously as curvature on the left-hand side and as stress-energy on the right.

Why is the vacuum energy 120 orders of magnitude smaller than quantum field theory predicts? Because quantum field theory evaluates the vacuum energy at $\psi = 0$, the symmetric maximum of the sombrero potential, not at the physical vacuum $\mathcal{V} \cong S^1$. The actual vacuum energy is the depth of the sombrero valley, which is finite and negative, set by the excess tension $\varepsilon(C)$ integrated over instantiated configurations.

Why do orbiting bodies orbit? Because the instantiated vacuum is topologically a circle and motion along that circle is the unique stable state. The geodesic equation of General Relativity correctly describes which paths are available; it cannot explain why any path is occupied. PST closes this gap: the Noether charge of the spontaneously broken U(1) symmetry of $\mathcal{V}$ is angular momentum, and it is conserved by construction, not by initial condition.

13.4. The unification of quantum mechanics and general relativity

The search for a unified theory of quantum mechanics and general relativity has been the central unsolved problem in theoretical physics for nearly a century. The difficulty has always been that each theory requires a different kind of background: quantum mechanics requires a fixed spacetime arena on which its field operators are defined; general relativity requires a dynamical spacetime that cannot be fixed without destroying the theory's content. Every attempt to quantize gravity by applying the Hilbert space formalism directly to the metric has encountered the problem of time, non-renormalisable divergences, and the absence of a sensible ground state. Every attempt to derive quantum mechanics from general relativity has encountered the measurement problem and the violation of locality.

PST resolves this not by finding a compromise between the two frameworks but by identifying the common origin from which both emerge. Quantum mechanics and general relativity are not competitors; they are projections of the same precausal structure at different distances from the modal threshold τ. The Hilbert space, the metric manifold, the uncertainty principle, and the geodesic equation are all derived objects, not primitives. They do not need to be made compatible because they were never incompatible at the foundational level; the apparent incompatibility arises only when each theory is taken as primitive and the two primitives are then compared. From the perspective of the precausal substrate, there is one structure, one threshold, one functor, and one projection. Quantum and classical are two names for two regimes of the same projection.

13.5. The quantitative standing of the theory

PST is not purely qualitative. Its central objects are mathematically explicit: the modal potential functional $\mathcal{F}(\psi, C)$ is a well-defined variational object; the threshold τ is defined by equation (16); the functor Φ is constructed in three explicit steps via the realization map ρ and the tension-to-metric assignment of equation (18); the projection operator Π is equation (19); the field equations follow as equation (25); the vacuum manifold radius is $r_0 = \sqrt{\varepsilon/(2b)}$; the Noether charge is equation (36). Each of these is a mathematical statement, not a metaphor.

The primary quantitative prediction of the theory is a correction to the Casimir effect. The projection operator Π introduces a spatial resolution scale d_0, the characteristic distance at which the realisation map ρ becomes sensitive to boundary geometry. At separations $d \sim d_0$, the PST projection diverges from the standard QFT result by a term scaling as d^{-6} rather than the standard d^{-4}, as given by equation (57). Section 10 derives d_0 from first principles via the dimensional reduction mechanism: modal sublimation compactifies three substrate dimensions at scale d_0, and the reduction formula $d_0^3 = \ell_*^5/\ell_P^2$ with the electroweak modal scale $M_* \approx 1.25$ TeV gives $d_0 \approx 7$ nm. At this value the Casimir correction at $d = 50$ nm is approximately 2%, detectable at next-generation precision; at $d = 20$ nm it reaches 12%, unambiguous against all known competing corrections. Existing measurements at 160–750 nm constrain $d_0 < 16$ nm, consistent with the prediction. A ratio test between measurements at two separations, equation (59), provides a clean, apparatus-independent signature distinguishable from finite-conductivity, surface roughness, and thermal corrections. This is a falsifiable prediction in the strict sense: a positive detection at $d_0 \approx 7$ nm would confirm the dimensional reduction mechanism; a null result at sub-50 nm precision at the 1% level would push M_* above 2 TeV and constrain the electroweak identification.

Beyond the Casimir prediction, PST makes a structural prediction about the scale dependence of spacetime dimensionality. Near the modal threshold, the geometry is high-dimensional; far from it, the settled vacuum selects four macroscopic dimensions. This scale-dependent dimensionality, with the spectral dimension decreasing from high values at short distances to four at large scales, is consistent with the independent result of Causal Dynamical Triangulations [32], providing cross-theory corroboration for the dimensional reduction mechanism.

13.6. The scope and limits of the theory

Intellectual honesty requires stating clearly what PST does not accomplish. The emergence chain reaches angular momentum and the topology of the vacuum as step 8. It does not yet derive the gauge group $SU(3) \times SU(2) \times U(1)$ of the Standard Model from first principles, though the $U(1)$ factor emerges directly from the vacuum manifold symmetry and the $SU(2)$ and $SU(3)$ factors are expected to arise from the higher-dimensional angular momentum structure near the threshold. The Born rule and the full Hilbert space formalism have not been derived from the structure of $\mathcal{F}$ and Π near threshold; the correspondence has been established qualitatively and deferred to future work. The value of d_0 has been derived conditionally in Section 10 under the electroweak condensate identification, giving $d_0 \approx 7$ nm; the unconditional derivation of the fundamental modal scale ℓ_* from (D, δ, T) alone, and the computation of the dimensionless coefficient ξ from the modal kernel, remain open. The canonical measure μ on configuration space was derived from first principles in Section 3; however, the gradient structure on $\mathcal{C}$ required by the functional calculus, and in particular the topology needed to define $\nabla_{\mathcal{C}}$, rests on structural choices whose full derivation from (D, δ) alone remains an open problem.

These are open questions, not weaknesses. They define the programme of work that PST opens up. A theory that identifies the correct primitive and constructs the correct emergence chain from it will inevitably leave specific derivations for subsequent work; that is the nature of a foundational theory. What matters at this stage is that the foundational structure is correct: that the primitive is genuinely irreducible, that the chain of entailments is logically tight, and that the resulting framework is consistent with everything established by existing physics while resolving the contradictions and open questions that existing physics cannot address from within itself.

13.7. On intelligibility and necessity

The deepest contribution of PST may not be its specific predictions but its answer to a question that physics has traditionally declined to answer: is the universe intelligible in its existence, or is its existence a brute fact beyond the reach of rational explanation?

The standard position in physics is that the most fundamental laws are simply given, that they describe the world without explaining why the world is described by them rather than by different laws, and that asking why the laws are as they are is a question that physics cannot and should not try to answer. This position is epistemically cautious but ontologically unsatisfying. It places the foundation of physical reality outside the reach of reason, not because reason has tried and failed but because the attempt has never been made within the framework of physics itself.

PST makes the attempt. It begins from a primitive that is not a law of nature but a logical necessity: the capacity for distinctions to exist is irreducible, and a domain with irreducible distinctions cannot remain tensionless, and a domain with tension above the modal threshold cannot remain uninstantiated. At every step the transition is not contingent but necessary. The result is not that the universe is as it is by law but that the universe is as it is by logic: the particular physics it displays is the necessary expression of the minimal ontological structure that any distinction-bearing domain must possess.

This is Leibniz's vision pursued with mathematical precision: sufficient reason, not as a metaphysical principle asserted from the armchair, but as a derivable consequence of the structure of the precausal substrate. The universe exists not because it was created or because it happens to be the case, but because the logical possibility of property differentiation is irreducible, and a substrate that cannot be otherwise is, in the only meaningful sense of the word, necessary.

In plain terms: the universe expresses structure in order to resolve its own intrinsic imbalances. This is not poetry. It is the theorem that follows from one primitive, through one threshold, by one functor, into one manifold. Everything else is projection.

13.8. Layman's Summary

Precausal Substrate Theory establishes that spacetime, causality, matter, the four fundamental forces, and the statistical uniformity of the cosmos all follow from a single, logically irreducible primitive: the capacity for distinctions to exist. No laws are postulated, no constants are assumed except as projective unit conversions, and no initial conditions are imposed. The theory makes a concrete experimental prediction (a modified Casimir force at nanometre separations) that distinguishes it from all current frameworks. The deepest implication is also the simplest: the universe exists not because something caused it, but because a reality in which no distinctions are possible is not a reality at all.

14. Open Problems

1. **Derivation of c_{LG} from the kernel K, and the coefficient ξ of the Casimir amplitude.**

   Section 10 (subsection "Simultaneous solution") has now carried out the algebraic solution of equations (A) and (B) simultaneously, using the Higgs identification $|a_{LG}| = m_h^2/4$. The result is the closed-form expression

$$M_* \;=\; m_h^{3/5}\, M_P^{2/5}\,(2\,c_{LG})^{-3/10}$$

(equation (73)). This expresses the fundamental modal scale M_* entirely in terms of the measured quantities m_h, M_P, and the gradient coefficient c_{LG} of the modal condensate. Using the LHC lower bound $M_* \approx 1.25$ TeV to invert, $c_{LG} \approx 1.0 \times 10^{19}$ (dimensionless, equation (75)).

The remaining open problem is the derivation of c_{LG} from first principles, rather than from inverting the LHC bracket. Equation (C) relates c_{LG} to G, v, and M_P via the gravitational coupling constraint, up to the normalization κ of the compact fibre projection integral. Once κ is computed from the kernel K, c_{LG} follows as a definite number, M_* becomes a parameter-free prediction, and Newton's constant G is derived unconditionally via equation (30) rather than taken as an input. The consistency check is that this derived c_{LG} must equal the value $\approx 1.0 \times 10^{19}$ already implied by the LHC bound; agreement would constitute a non-trivial first-principles derivation of the electroweak hierarchy.

Separately, the dimensionless coefficient ξ in the Casimir correction (equation (57)) requires computing the modal density of states from the same kernel K. Both quantities, c_{LG} and ξ, reduce to integrals over the kernel geometry; computing them is the single calculational task that closes the remaining parametric freedom in the theory's predictions.

2. **Derivation of QM and QFT from PST.** The threshold regime was identified as the precausal origin of quantum indeterminacy and the projected field theory was constructed in Sections 7 and 8. Three open items remain. *(a)* A complete derivation of the Hilbert space formalism, Born rule, and measurement postulate from the structure of $\mathcal{F}$ and Π near threshold. *(b)* The fermionic sector. Section 8 establishes that the spinor bundle exists on (M, g) without topological obstruction, and sharpens the open problem to the following: whether the Bernoulli measure μ gives non-zero weight to path-integral sectors transforming in the spinor representation $j = 1/2$ of the threshold SU(2), and whether Π projects those sectors to Dirac fermions on M. The number of generations $N_{\text{gen}} = 3$, the Yukawa hierarchy, and the CKM matrix are contingent on the resolution of this problem. *(c)* The full interacting vacuum and its renormalisation at the d_0-regulated scale.

3. **Connection to the Standard Model.** Section 8 derives the full gauge group $G_{\text{SM}} = \text{SU}(3)_c \times \text{SU}(2)_L \times \text{U}(1)_Y$ as the automorphism group of the projected substrate vacuum: $\text{U}(1)_Y$ from the vacuum-manifold phase rotation, $\text{SU}(2)_L$ from the $S^3 \cong \text{SU}(2)$ threshold topology, and $\text{SU}(3)_c$ from the $N = 3$ compact fibre. Section 8 further establishes gauge anomaly cancellation as a structural theorem. The remaining open problems in the Standard Model connection are the fermionic sector (Problem 2b above), the number of generations $N_{\text{gen}} = 3$, the Yukawa coupling hierarchy, and the Cabibbo–Kobayashi–Maskawa mixing matrix.

4. **Grounding the gradient structure on configuration space.** The gradient term $|\nabla_{\mathcal{C}}\psi|^2$ in the modal potential functional presupposes a topology on $\mathcal{C}$ that is not entailed by (D, δ) alone. The natural symmetric-difference metric on $\mathcal{P}(D)$ is a structurally motivated candidate, but a rigorous derivation of the differential structure required for functional calculus from the primitive distinction relation, without importing geometric or topological assumptions, remains to be given.

5. **Explicit derivation of macroscopic orbital angular momentum via Π.** Section 7 establishes that the broken U(1) symmetry of the vacuum manifold generates a conserved Noether charge $L = 2r_0^2\dot{\theta}$ in configuration space. The structural argument for identifying this charge with physical orbital angular momentum via the projection Π is clear. However, the explicit computation showing how L maps through Π to the orbital angular momentum of a gravitationally bound body moving on a geodesic, connecting the configuration-space Noether charge to the classical angular momentum observable, is a derivation that remains to be carried out in full.

14.1. Problems resolved within this paper

Four of the seven open problems listed in earlier drafts have been resolved and removed from this section. The measure μ on configuration space was derived in Section 3 from two invariance requirements alone, automorphism-invariance and complementation invariance, establishing it uniquely as the Bernoulli product measure $\mu = \bigotimes_{a\in D} \mathrm{Bern}(1/2)$, with no geometric or probabilistic primitive assumed. The uniqueness of Φ was established in Section 4: the functor is unique up to diffeomorphism, its ambiguity is completely and exactly characterised by diffeomorphism invariance, and the Lorentzian signature $(-,+,+,+)$ is uniquely fixed by the single extremal direction of asymmetric tension. The derivation of G from Π was addressed in Section 7: the gravitational constant is not independent of d_0 but equals $\alpha_G \cdot d_0^2$ in natural units, reducing the question to determining d_0. The conditional derivation of d_0 was completed in Section 10: the dimensional reduction formula $d_0^3 = \ell_*^5/\ell_P^2$ with $N = 3$ compact substrate dimensions and $M_* \approx 1.25$ TeV gives $d_0 \approx 7$ nm, establishing the Casimir correction as a predicted detectable signal. Two new open problems have been added in this version in response to critical scrutiny of the formalism: the grounding of the gradient structure on configuration space (Problem 4), and the explicit derivation of macroscopic orbital angular momentum from the Noether charge via Π (Problem 5). These additions reflect the programme of honest self-examination that foundational work requires.

Precausal Substrate Theory, Ralf Meelker, 2026.

14.2. Layman's Summary

PST derives the bosonic sector of the Standard Model (gauge bosons, the Higgs field, and gravitons) but the fermionic sector (quarks, leptons, and the spin-statistics connection) remains the central open problem: the mechanism by which the substrate generates half-integer spin has not yet been constructed. Further open questions include why there are exactly three families of quarks and leptons, why the top quark is so much heavier than the electron, and what determines the mixing angles between generations. Resolving these would transform PST from a framework that derives the *structure* of particle physics into one that derives its *content*.

Appendix: For the General Reader

What the theory is trying to do

Physics has been extraordinarily successful at describing how the universe works. General Relativity tells us how gravity curves space and time. Quantum Mechanics tells us how subatomic particles behave. Both theories work to extraordinary precision in their respective domains. But neither theory asks the question that comes one step earlier: *why is there a universe to describe at all?* Why does space exist? Why does time have a direction? Why does matter exist and why does it weigh something? The standard answer in physics is: that is just how it is; we describe what we find and do not ask for deeper reasons.

Precausal Substrate Theory (PST) refuses that answer. It holds that those questions do have answers, and that the answers follow logically from a single, irreducible primitive, one that does not itself presuppose space, time, or any physical structure.

The single primitive

The entire theory is built from one observation: *things can be different from each other*. That is it. Before there is space, before there is time, before there is matter or energy or any physical law, there must be the possibility of distinction: the most minimal fact that two things are not the same.

Call this **property differentiation**. It cannot be explained in terms of anything simpler, because any explanation would already need two concepts to be different from each other. It is the floor of all reasoning.

The philosopher Leibniz saw something like this: without sufficient reason, nothing would differ from anything else and the universe would be perfectly featureless, which is to say, there would be no universe. The logician G. Spencer-Brown made the same move in formal logic: the entirety of Boolean algebra, the foundation of all digital computing, follows from a single instruction, “draw a distinction.” PST takes that insight and asks: what does it imply, physically, if you follow it all the way?

The eight steps

Once distinction is possible, a chain of logical consequences follows, each step unavoidable given the one before:

1. **Distinctions exist.** Some things differ from others. This is the single primitive.
2. **Asymmetric tension.** Any collection of distinct things carries an imbalance: it is not in perfect equilibrium with its mirror image. Think of a handful of different-coloured marbles: taken as a group, they are structured in a way that its “opposite” group (the marbles not chosen) is not. PST calls this imbalance asymmetric tension.
3. **A breaking point.** As the imbalance grows, there comes a level at which it can no

longer be "contained" in an abstract, pre-physical state. A threshold is crossed.

4. **Modal sublimation.** When the threshold is crossed, the abstract collection of distinctions cannot remain abstract. It is *compelled*, not by any force, not by any event, but by logical necessity, to become concrete. PST calls this modal sublimation. The universe does not arise from a bang at a single point in time; it arises as a logical inevitability, without a temporal or spatial origin.

5. **Space and time appear together.** The act of becoming concrete simultaneously creates space and time. There is no moment before spacetime at which the universe began, because "before" requires time, and time is part of what is being created. Asking "what happened before the Big Bang?" is like asking what is north of the North Pole.

6. **Causality.** With spacetime comes the structure of cause and effect: the fact that some events can influence others in one direction only. Causality is not a separate law of nature; it is a property of the spacetime that modal sublimation produces.

7. **Matter and energy.** The original asymmetric tension, now projected into spacetime, shows up as what we call matter and energy. The marble and the wood of Einstein's field equations (geometry and matter) are both expressions of the same original imbalance.

8. **Angular momentum.** The mathematical structure of the vacuum, the ground state of the universe after sublimation, has the shape of a circle in an internal sense. From that circular shape, by a theorem about conserved quantities (Noether's theorem), comes angular momentum: the tendency of everything from electrons to planets to orbit rather than to stand still.

The experimental test

The theory makes a concrete prediction distinguishable from standard quantum physics. The Casimir effect is the tiny attractive force between two uncharged metal plates placed extremely close together, a force that exists because the vacuum itself has energy. Standard physics predicts how this force scales with distance. PST predicts a small additional correction scaling differently with distance (as d^{-6} rather than d^{-4}), arising from the fact that the vacuum has a finite coherence length: the compactification radius d_0 of the substrate's compact dimensions. From first principles (Section 10), PST derives $d_0 \approx 7$ nm. At plate separations near 50 nm the correction reaches approximately 2%, within reach of next-generation precision experiments; at 20 nm it reaches 12%, unambiguous against all competing effects. If that correction is found at the predicted scale, it would confirm that the vacuum has the pre-geometric discrete structure PST describes, and simultaneously fix the fundamental modal scale at approximately 1.25 TeV, a value accessible to particle accelerators.

What makes this different

Every other physical theory assumes its arena. General Relativity assumes that spacetime exists and asks how it curves. Quantum Mechanics assumes that a mathematical space of states exists and asks how things evolve in it. String Theory assumes that strings propagate through a background spacetime. None of them asks why that arena must exist rather than nothing.

PST operates beneath every such assumption. It does not assume spacetime, matter, causality, or any of the structures that other theories take for granted. It derives them, one by one, from the single fact that things can differ from each other. The result is not a replacement for existing physics, General Relativity and Quantum Mechanics remain correct within their domains, but a foundation beneath them, an explanation of why they work and why the universe they describe had to exist at all.

The deepest philosophical claim is also the simplest: the universe exists not because something caused it, but because the alternative, a reality in which no distinctions are possible, is not a reality at all.

References

[1] A. Einstein, “Die Feldgleichungen der Gravitation,” *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, 844 (1915). [archive scan]

[2] R. Penrose, “Gravitational collapse and space-time singularities,” *Phys. Rev. Lett.* **14**, 57 (1965). doi:10.1103/PhysRevLett.14.57

[3] S. W. Hawking and R. Penrose, “The singularities of gravitational collapse and cosmology,” *Proc. R. Soc. Lond. A* **314**, 529 (1970). doi:10.1098/rspa.1970.0021

[4] L. Bombelli, J. Lee, D. Meyer, R. D. Sorkin, “Space-time as a causal set,” *Phys. Rev. Lett.* **59**, 521 (1987). doi:10.1103/PhysRevLett.59.521

[5] R. D. Sorkin, “Forks in the road, on the way to quantum gravity,” *Int. J. Theor. Phys.* **36**, 2759 (1997). arXiv:gr-qc/9706002

[6] C. Rovelli and L. Smolin, “Loop space representation of quantum general relativity,” *Nucl. Phys. B* **331**, 80 (1990). doi:10.1016/0550-3213(90)90019-A

[7] A. Ashtekar, “New variables for classical and quantum gravity,” *Phys. Rev. Lett.* **57**, 2244 (1986). doi:10.1103/PhysRevLett.57.2244

[8] V. L. Ginzburg and L. D. Landau, “On the theory of superconductivity,” *Zh. Eksp. Teor. Fiz.* **20**, 1064 (1950). Reprinted in L. D. Landau, *Collected Papers*, Pergamon Press (1965). INSPIRE-HEP:9135

[9] S. Mac Lane, *Categories for the Working Mathematician*, Springer, New York (1971). doi:10.1007/978-1-4612-9839-7

[10] C. W. Misner, K. S. Thorne, J. A. Wheeler, *Gravitation*, W. H. Freeman, San Francisco (1973). Princeton UP reprint

[11] H. B. G. Casimir, "On the attraction between two perfectly conducting plates," *Proc. Kon. Ned. Akad. Wetensch.* **51**, 793 (1948). [original PDF]

[12] S. K. Lamoreaux, "Demonstration of the Casimir force in the 0.6 to 6 μm range," *Phys. Rev. Lett.* **78**, 5 (1997). doi:10.1103/PhysRevLett.78.5

[13] G. Bressi, G. Carugno, R. Onofrio, G. Ruoso, "Measurement of the Casimir force between parallel metallic surfaces," *Phys. Rev. Lett.* **88**, 041804 (2002). arXiv:quant-ph/0203002

[14] R. S. Decca, D. López, E. Fischbach, D. E. Krause, "Measurement of the Casimir force between dissimilar metals," *Phys. Rev. Lett.* **91**, 050402 (2003). arXiv:quant-ph/0306136

[15] M. Bordag, U. Mohideen, V. M. Mostepanenko, "New developments in the Casimir effect," *Phys. Rept.* **353**, 1 (2001). arXiv:quant-ph/0106045

[16] N. Arkani-Hamed, S. Dimopoulos, G. Dvali, "The hierarchy problem and new dimensions at a millimeter," *Phys. Lett. B* **429**, 263–272 (1998). arXiv:hep-ph/9803315

[17] E. P. Verlinde, "On the origin of gravity and the laws of Newton," *JHEP* **04** (2011) 029. arXiv:1001.0785

[18] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Phys. Rev. Lett.* **75**, 1260 (1995). arXiv:gr-qc/9504004

[19] A. Einstein, "Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie," *Sitzungsber. Preuss. Akad. Wiss.* **142** (1917). ADS:1917SPAW.......142E

[20] E. Noether, "Invariante Variationsprobleme," *Nachr. D. König. Gesellsch. D. Wiss. Zu Göttingen, Math-phys. Klasse*, 235 (1918). arXiv:physics/0503066

[21] D. Lovelock, "The Einstein tensor and its generalizations," *J. Math. Phys.* **12**, 498–501 (1971). doi:10.1063/1.1665613

[22] J. A. Wheeler and E. F. Taylor, *Spacetime Physics*, 2nd ed., W. H. Freeman, New York (1992). eftaylor.com/spacetimephysics

[23] P. A. M. Dirac, *The Principles of Quantum Mechanics*, Clarendon Press, Oxford (1930); 4th ed. (1958). [archive scan]

[24] J. von Neumann, *Mathematische Grundlagen der Quantenmechanik*, Springer, Berlin (1932). English translation: *Mathematical Foundations of Quantum Mechanics*, Princeton University Press (1955; reprint 2018). Princeton UP

[25] R. P. Feynman and A. R. Hibbs, *Quantum Mechanics and Path Integrals*, McGraw-Hill, New York (1965). Dover emended edition (2010). Dover Publications

[26] M. B. Green, J. H. Schwarz, E. Witten, *Superstring Theory* (2 vols.), Cambridge University Press (1987). Cambridge UP

[27] E. Witten, "String theory dynamics in various dimensions," *Nuclear Physics B* **443** (1995) 85–126. doi:10.1016/0550-3213(95)00158-O

[28] T. Kaluza, "Zum Unitätsproblem der Physik," *Sitzungsber. Preuss. Akad. Wiss. Berlin* (1921) 966–972. ADS:1921SPAW.......966K

[29] O. Klein, "Quantentheorie und fünfdimensionale Relativitätstheorie," *Zeitschrift für Physik* **37** (1926) 895–906. doi:10.1007/BF01397481

[30] N. Arkani-Hamed, S. Dimopoulos, G. Dvali, "The hierarchy problem and new dimensions at a millimeter," *Physics Letters B* **429** (1998) 263–272. doi:10.1016/S0370-2693(98)00466-3

[31] L. Randall, R. Sundrum, "A large mass hierarchy from a small extra dimension," *Physical Review Letters* **83** (1999) 3370–3373. doi:10.1103/PhysRevLett.83.3370

[32] J. Ambjørn, J. Jurkiewicz, R. Loll, "Reconstructing the universe," *Physical Review D* **72** (2005) 064014. doi:10.1103/PhysRevD.72.064014

[33] G. 't Hooft, "Dimensional reduction in quantum gravity," arXiv:gr-qc/9310026 (1993). arXiv:gr-qc/9310026

[34] J. Maldacena, "The large N limit of superconformal field theories and supergravity," *International Journal of Theoretical Physics* **38** (1999) 1113–1133. doi:10.1023/A:1026654312961

[35] I. Kant, *Kritik der reinen Vernunft*. Hartknoch, Riga, 1781. English translation: *Critique of Pure Reason*, trans. N. Kemp Smith, Macmillan, London, 1929. [archive scan, Kemp Smith trans.]

[36] J. Henson, "The causal set approach to quantum gravity," in D. Oriti (ed.), *Approaches to Quantum Gravity*, Cambridge University Press, Cambridge, 2009, pp. 393–413. doi:10.1017/CBO9780511575549.025

[37] J. B. Hartle and S. W. Hawking, "Wave function of the universe," *Physical Review D* **28** (1983) 2960–2975. doi:10.1103/PhysRevD.28.2960

[38] R. Geroch, "Domain of dependence," *Journal of Mathematical Physics* **11** (1970) 437–449. doi:10.1063/1.1665157

[39] R. M. Wald, *General Relativity*. University of Chicago Press, Chicago, 1984. University of Chicago Press

[40] G. Ryle, *The Concept of Mind*. Hutchinson, London, 1949. [archive scan]

[41] H. Putnam, *Reason, Truth and History*. Cambridge University Press, Cambridge, 1981. doi:10.1017/CBO9780511625398

[42] G. W. Leibniz, "La Monadologie," 1714. English translation: *The Monadology*, in *Philosophical Papers and Letters*, 2nd ed., trans. L. E. Loemker, Reidel, Dordrecht, 1969. [archive scan, original French]

[43] A. Vilenkin, "Creation of universes from nothing," *Physics Letters B* **117** (1982) 25–28. doi:10.1016/0370-2693(82)90866-8

[44] J. Goldstone and J. C. Ward, "Field theories with superconductor solutions," *Il Nuovo*

Cimento **19** (1961) 154–164. doi:10.1007/BF02812722

[45] P. A. M. Dirac, "The cosmological constants," *Nature* **139** (1937) 323. doi:10.1038/139323a0

[46] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "Observation of gravitational waves from a binary black hole merger," *Physical Review Letters* **116** (2016) 061102. doi:10.1103/PhysRevLett.116.061102

[47] J. D. Bekenstein, "Black holes and entropy," *Physical Review D* **7** (1973) 2333–2346. doi:10.1103/PhysRevD.7.2333

[48] S. W. Hawking, "Particle creation by black holes," *Communications in Mathematical Physics* **43** (1975) 199–220. doi:10.1007/BF02345020

[49] A. Connes, *Noncommutative Geometry*. Academic Press, San Diego, 1994. [free PDF, author's website]

[50] A. H. Guth, "Inflationary universe: A possible solution to the horizon and flatness problems," *Physical Review D* **23** (1981) 347–356. doi:10.1103/PhysRevD.23.347

[51] A. D. Linde, "Chaotic inflation," *Physics Letters B* **129** (1983) 177–181. doi:10.1016/0370-2693(83)90837-7

[52] B. S. DeWitt, "Quantum Theory of Gravity. I. The Canonical Theory," *Physical Review* **160** (1967) 1113–1148. doi:10.1103/PhysRev.160.1113

[53] G. Spencer-Brown, *Laws of Form*. Allen & Unwin, London, 1969. [archive.org scan]

[54] J. A. Wheeler, "Information, Physics, Quantum: The Search for Links," in W. H. Zurek (ed.), *Complexity, Entropy and the Physics of Information*, Addison-Wesley, Redwood City, CA, 1990. doi:10.1201/9780429500459-19

[55] M. Tegmark, "The Mathematical Universe," *Foundations of Physics* **38** (2008) 101–150. doi:10.1007/s10701-007-9186-9 [arXiv:0704.0646]

[56] J. Ladyman, "What is Structural Realism?" *Studies in History and Philosophy of Science* **29** (1998) 409–424. doi:10.1016/S0039-3681(98)80129-5

[57] S. French, *The Structure of the World: Metaphysics and Representation*. Oxford University Press, Oxford, 2014. ISBN 978-0-19-968484-7.

[58] D. Bohm, *Wholeness and the Implicate Order*. Routledge & Kegan Paul, London, 1980. [archive.org scan]

[59] G. W. F. Hegel, *Science of Logic*, trans. A. V. Miller. Allen & Unwin, London, 1969 (original: *Wissenschaft der Logik*, Cotta, Nuremberg, 1812–1816). [complete text, Marxists Internet Archive]

[60] A. Friedmann, "Über die Krümmung des Raumes," *Zeitschrift für Physik* **10** (1922) 377–386. doi:10.1007/BF01332580

[61] A. G. Riess *et al.* (High-z Supernova Search Team), "Observational evidence from supernovae for an accelerating universe and a cosmological constant," *The Astronomical Journal* **116** (1998) 1009–1038. doi:10.1086/300499 [arXiv:astro-ph/9805200]

[62] S. Perlmutter *et al.* (Supernova Cosmology Project), "Measurements of Ω and Λ from 42 high-redshift supernovae," *The Astrophysical Journal* **517** (1999) 565–586. doi:10.1086/307221 [arXiv:astro-ph/9812133]

[63] N. Aghanim *et al.* (Planck Collaboration), "Planck 2018 results. VI. Cosmological parameters," *Astronomy & Astrophysics* **641** (2020) A6. doi:10.1051/0004-6361/201833910 [arXiv:1807.06209]

[64] H. Bondi and T. Gold, "The steady-state theory of the expanding universe," *Monthly Notices of the Royal Astronomical Society* **108** (1948) 252–270. doi:10.1093/mnras/108.3.252

[65] F. Hoyle, "A new model for the expanding universe," *Monthly Notices of the Royal Astronomical Society* **108** (1948) 372–382. doi:10.1093/mnras/108.5.372

[66] A. A. Penzias and R. W. Wilson, "A measurement of excess antenna temperature at 4080 Mc/s," *The Astrophysical Journal* **142** (1965) 419–421. doi:10.1086/148307