

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 spacetime, causality, matter, and the fundamental forces are not themselves fundamental but arise from logic: a precausal substrate whose sole primitive is **property differentiation**, the capacity for distinctions to exist, with directional content in the 7-dimensional direction algebra V_7 forced by Hurwitz’s classification of normed division algebras. Property differentiation carries an asymmetric vector tension; past a modal threshold, modal sublimation produces instantiated geometry along the substrate’s net direction, with causality born together with geometry. The Big Bang is a perspectival artifact rather than a first event. 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.

The Core Mathematics. The substrate is a real Connes spectral triple of KO-dimension six with sign pattern $(\varepsilon, \varepsilon', \varepsilon'') = (+, +, -)$, Connes’ Lorentzian-compatible value (Computation 3). KO-dimension is an algebraic mod-8 invariant of real K-theory (“KO” = Klasse Orthogonal), not a count of physical dimensions. The substrate’s Mosco limit yields spacetime $M = \mathbb{R} \times S^3$ with KO-dimension four (Computation 4); the product $M \times F$ has total **KO-dimension ten** $\equiv 2 \pmod{8}$, the Connes Standard-Model value. Spacetime dimension $n = 4$ is forced. The substrate’s internal algebra is derived as $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, generating the gauge group $SU(3) \times SU(2) \times U(1)$ of the strong, weak, and electromagnetic forces. Computation 19 verifies that the parity grading γ_K equals the $Cl(0, 6)$ chirality grading $i\omega$ exactly (Jordan-Wigner, machine precision), so the chirality projector is Furey’s primitive idempotent [102] and Furey’s calculation delivers A_F as the matter-representation commutant (§8.8). No Chamseddine-Connes-Marcolli classification is invoked.

The Unified Physics. The same spectral-triple structure delivers the four otherwise-disjoint foundations of contemporary physics. *General relativity:* the Einstein-Hilbert action emerges as the leading term of the spectral action. *Quantum mechanics:* canonical commutation relations arise via Mosco convergence of Boolean Dirichlet forms. *The Standard Model:* gauge bosons and the Higgs appear as inner fluctuations of the Dirac operator. *Matter:* Boolean $\{0, 1\}$ distinctions are fermionic occupation numbers via the CAR algebra, making spin-statistics structural. The structural derivation of $N_{\text{gen}} = 3$ generations follows from the asymmetric-tension postulate (§1.7): the substrate’s net directional bias $\hat{\tau} = T(D)/|T(D)| \in S^6 \subset V_7$ selects three quaternionic sub-algebras of \mathbb{O} containing $\hat{\tau}$, each carrying one generation’s Standard-Model matter content (§8.17, Computations 46–48); the dimension 7 of V_7 is forced by Hurwitz’s classification of normed division algebras over \mathbb{R} .

Predictions. PST predicts a d^{-6} Casimir correction at nanometre separations with

parameter-free coefficient $\xi = 90/\pi^2$ at the substrate coherence length $d_0 < 16$ nm. The electroweak modal scale is $M_* = 4\pi m_h \approx 1.573$ TeV, structurally consistent with electroweak precision: $\Delta S = \Delta T = 0$ at M_* (§10.9). The cosmological-constant sign is delivered structurally (§7.14); its magnitude, the substrate coherence scale d_0 , the Yukawa hierarchies, and CKM mixing are contingent and not pinned by the postulates alone.

Open piece. Two analytical items remain in the rigorous-proof tail of the Dirac-convergence step (§14.2). The single-mode L_round closure is already a theorem (Computation 50, via $D_{\text{sub}}^2 = D \cdot I$), and the block-decomposition identity reduces the tail-sum case to an explicit operator-norm bound on the off-diagonal block of the Walsh tail in the D_{sub} eigenbasis (Computation 51). The remaining open work is (i) the closed-form proof of that off-diagonal bound and (ii) Berezin-transform compatibility plus spectral-action invariance for the refined symmetric-monomial bridge that closes L_comm (Computations 39–42).

Contents

1	Introduction	8
1.1	The incompatibility of quantum mechanics and general relativity	8
1.2	Assumed backgrounds and the missing question	8
1.3	The precausal starting point	9
1.4	The foundational object: a spectral triple of total KO-dimension ten	9
1.5	Principal consequences	11
1.6	Mathematical formulation and predictions	12
1.7	The foundational postulates of PST	13
1.8	Structure of this paper	14
1.9	Layman’s Summary	15
2	Foundational Ontology	15
2.1	The case for ontological minimality	15
2.2	Property differentiation	16
2.3	Mathematical formulation	17
2.4	Asymmetric tension	18
2.5	Layman’s Summary	21
3	The Modal Threshold	21
3.1	Why a modal threshold must exist	21
3.2	Boundary condition	23
3.3	The order parameter and its ontological status	24
3.4	Variational formulation of τ	25
3.5	Stability analysis and the two regimes	30
3.6	The bifurcation point and the derivation of τ	32
3.7	Layman’s Summary	33
4	Modal Sublimation	34
4.1	The character of the transition	34
4.2	The nontemporal character	34
4.3	The nonmechanistic character	35
4.4	The nonenergetic character	35
4.5	The strictly modal character	36
4.6	Categorical formulation	36
4.7	Explicit construction of Φ	38
4.8	Uniqueness of Φ and the origin of diffeomorphism invariance	40
4.9	Layman’s Summary	41
5	Coinstantiation of Spacetime and Causality	42
5.1	Causality as coinstantiated structure	42

5.2	Mathematical formulation of coinstantiation	42
5.3	The direction of time as a structural consequence	43
5.4	The causal order and the degenerate vacuum	44
5.5	Category error and the dissolution of first-cause questions	44
5.6	Contrast with competing approaches	45
5.7	Layman’s Summary	46
6	The Emergence Chain	46
6.1	Overview and the eight-step summary table	46
6.2	The character of the chain	47
6.3	Step 1 to Step 2: from differentiation to tension	48
6.4	Step 2 to Step 3: from tension to threshold	48
6.5	Step 3 to Step 4: from threshold to sublimation	49
6.6	Step 4 to Steps 5 and 6: from sublimation to spacetime and causality	49
6.7	Steps 6 and 7: from causality and geometry to matter	50
6.8	Step 8: from geometry and the vacuum to angular momentum	50
6.9	The chain as a closed logical structure	51
6.10	Layman’s Summary	52
7	Energy, Matter, and Fields	52
7.1	The source problem in General Relativity	52
7.2	The stress-energy tensor as projected tension	53
7.3	The field equations as a conservation statement	54
7.4	The F–S dichotomy	54
7.5	The projected action and the emergence of the Einstein–Hilbert term	55
7.6	The Seeley–DeWitt argument for the functional form of equation (37)	56
7.7	Dirac-operator convergence: structural prerequisites and the residual analytic estimate	60
7.8	Kernel normalisation: the finite-internal-factor trace	67
7.9	Newton’s constant as a consistency relation	68
7.10	Uniqueness of $G_{\mu\nu}$: diffeomorphism invariance and Lovelock’s theorem	68
7.11	The equivalence principle as a theorem of single-source projection	69
7.12	The gravitational constant as a unit conversion	70
7.13	Derivation of G from the projection scale	70
7.14	The cosmological constant: substrate energy versus emergent vacuum density	71
7.15	The Degenerate Vacuum Manifold and the Origin of Angular Momentum	73
7.16	The incomplete vacuum of standard physics	74
7.17	Angular momentum as Noether charge	75
7.18	Angular momentum tensor and the Goldstone–geodesic connection	76
7.19	The projected field theory: emergence of quantum fields	78
7.20	Interacting vacuum and d_0 -regulated renormalisation	81
7.21	Layman’s Summary	84

8	Gauge Structure, Quantization, and the Standard Model	84
8.1	From $U(1)$ to the full Standard Model gauge group	84
8.2	Path-integral quantization of the projected theory	85
8.3	From global to local gauge symmetry: the photon as a connection	86
8.4	$SU(2)_L$ from the threshold topology: the electroweak doublet	86
8.5	$SU(3)_c$ as the colour factor of the internal algebra	87
8.6	Octonionic $SU(3)$ inside $Cl(0, 6)$: rigorous Casimir decomposition	87
8.7	Consistency of the octonionic interpretation across PST’s structural results	89
8.8	Closure of Conjecture 1: the algebra identity from P1–P3 via Computation 3 and Furey 2014	90
8.9	The Standard Model gauge group as substrate automorphisms	92
8.10	Gauge anomaly cancellation as a structural theorem	93
8.11	$U(1)_Y$ hypercharges from PST structure	94
8.12	The Standard Model Lagrangian as leading-order projection	95
8.13	The fermionic sector: structure and the sharpened open problem	96
8.14	Fermionic excitations from the Boolean structure: the CAR algebra identification	96
8.15	$SU(3)_c$ from the internal algebra	98
8.16	Fermionic representation content under Π	98
8.17	The three-generation count from the V_7 directional content	100
8.18	Yukawa hierarchy and CKM mixing as contingent in $T(C)$: a structural no-go	102
8.19	Chirality of $SU(2)_L$ as a structural consequence	102
8.20	Layman’s Summary	103
9	Physical Predictions	103
9.1	Structural and specific predictions	103
9.2	The Born rule from the Bernoulli measure	104
9.3	Canonical commutation relations from the Boolean gradient	105
9.4	Quantum Fluctuations	106
9.5	Spacetime Curvature	107
9.6	The weak-field limit and the recovery of Newtonian gravity	108
9.7	Stress-Energy as Projected Tension	109
9.8	Vacuum Phenomena and the Casimir Effect	109
9.9	The parameter-free exponent and the structure of falsifiability	110
9.10	Observational constraint	112
9.11	Layman’s Summary	113
10	Grounding d_0: Modal Coherence and the Casimir Prediction	113
10.1	The default assumption that must be dropped	113
10.2	Two distinct scales: microscopic grain versus coherence length	113
10.3	Modal sublimation and the substrate scales	114
10.4	The fundamental modal scale and the electroweak identification	115
10.5	The substrate coherence scale	116

10.6	Casimir detectability at each separation regime	116
10.7	Falsifiability and the experimental programme	117
10.8	Experimental status of M_* and the ADD comparison	118
10.9	Wilson coefficients and electroweak precision: $\Delta S = \Delta T = 0$	118
10.10	Correction budget for the Casimir measurement	119
10.11	Layman’s Summary	120
11	Cosmology: Expansion and Acceleration from Ongoing Sublimation	121
11.1	PST’s reframing of the standard cosmological picture	121
11.2	The Cosmological Principle as a theorem of the Bernoulli measure	121
11.3	Horizon and flatness problems dissolved	122
11.4	Ongoing instantiation: why the configuration space does not dilute	122
11.5	Equation of state and the sign of the cosmological constant	123
11.6	Friedmann equations and cosmic acceleration	123
11.7	The rate problem	124
11.8	Observational commitments	124
11.9	Comparison with steady-state cosmology	125
11.10	Layman’s Summary	126
12	Relation to Existing Formalisms	126
12.1	PST as foundation, not competitor	126
12.2	General Relativity	126
12.3	Quantum Mechanics and Quantum Field Theory	127
12.4	Thermodynamics and Black Hole Entropy	128
12.5	Causal Set Theory	128
12.6	Loop Quantum Gravity	129
12.7	String Theory and M-Theory	130
12.8	Kaluza-Klein Theory and Extra Dimensions	130
12.9	Large Extra Dimensions and Braneworld Models	131
12.10	Causal Dynamical Triangulations	132
12.11	The Holographic Principle and AdS/CFT	132
12.12	Non-Commutative Geometry	133
12.13	Inflationary Cosmology and the Multiverse	133
12.14	Spencer-Brown and Distinction-Based Ontologies	134
12.15	Wheeler’s “It from Bit” and the Mathematical Universe	135
12.16	Structural Realism and Bohm’s Implicate Order	136
12.17	Layman’s Summary	136
13	Conclusion	137
13.1	The shape of the argument	137
13.2	The central claim, stated precisely	137
13.3	What the emergence chain establishes	138

13.4	The open questions PST closes	139
13.5	The unification of quantum mechanics and general relativity	140
13.6	The quantitative standing of the theory	140
13.7	The scope and limits of the theory	141
13.8	On intelligibility and necessity	142
13.9	Layman's Summary	142
14	Summary of Results	143
14.1	Research problems resolved	143
14.2	Open research work	145
14.3	Layman's Summary	146
Appendix A: Notation Conventions		148
Appendix B: Bibliographic References		150
Appendix C: Computation Scripts		157

Website	https://meelker.nl
Full Paper (PDF)	https://meelker.nl/pst-paper-v19.13.pdf
Computations	https://meelker.nl/computations
LaTeX Source	https://meelker.nl/pst-paper.tex
Video (animation 1)	https://meelker.nl/video

Animations:

Modal sublimation into instantiated geometry

<https://meelker.nl/pst-animation1>

The eight-step chain of logical entailment from distinction to orbital motion

<https://meelker.nl/pst-animation2>

Asymmetric tension crossing the modal threshold

<https://meelker.nl/pst-animation3>

The Casimir discreteness correction

<https://meelker.nl/pst-animation4>

Vacuum topology and the origin of orbital motion

<https://meelker.nl/pst-animation5>

The pitchfork: spontaneous symmetry breaking at the modal threshold

<https://meelker.nl/pst-animation6>

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, relations, 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. The foundational object: a spectral triple of total KO-dimension ten

The substrate does not, on its own, single out a number of dimensions. What it provides is an algebra of distinctions with an intrinsic involution (the complementation $C \mapsto \bar{C}$) and a magnitude (the tension norm). The question that fixes the structure is not how many

dimensions the substrate has, but what the substrate must become for a Lorentzian quantum field theory to be definable on it. That requirement, the existence of a well-defined fermionic action continuing from the Euclidean to the Lorentzian, is a demand of consistency, and it is the same demand that the incompatibility of quantum mechanics and general relativity, with which this section opened, leaves unmet by every background-assuming theory.

Connes' noncommutative geometry answers that demand with a single integer, the KO-dimension, defined modulo eight by Bott periodicity. It is not the dimension of any manifold; it is a property of the real structure J and the grading γ , fixing the signs of $J^2 = \pm 1$, $J\gamma = \pm\gamma J$, and $JD = \pm DJ$. Of the eight values, the canonical one for a product of spacetime with an internal space, the value at which the Euclidean fermionic action carries the sign that continues to a well-defined Lorentzian action, is KO-dimension two.

PST meets this value structurally. The substrate's real structure has KO-dimension six (a computed result, holding for an odd number of distinction bits; Section 14.2). The Mosco limit of the substrate's Dirichlet form produces a four-dimensional Lorentzian spacetime, of KO-dimension four. KO-dimension is additive under the product, so the physical object, spacetime tensored with the substrate as its internal factor, has total KO-dimension $4 + 6 = 10 \equiv 2 \pmod{8}$. Among all totals that land on the Lorentzian-compatible value two, ten is the smallest: any larger total would require a larger spacetime (which the Mosco limit does not produce) or internal structure beyond the substrate. The foundational number of PST is therefore this total KO-dimension ten, from which the four of spacetime and the six of the internal factor are the *minimal split*: $4 + 6 = 10 \equiv 2 \pmod{8}$ is the lowest sum with both summands non-negative and the total congruent to two modulo eight (the Connes Lorentzian value), given the verified parity-selected substrate KO-6 (Computation 3) and the Mosco-limit spacetime KO-4 (Computation 4). "Minimal" is therefore conditional on the parity selection but does not require an extra spectral-triple axiom beyond what Computations 3 and 4 verify.

This total has nothing to do with the ten of superstring theory. There, ten is the dimension of a smooth target manifold, fixed by cancellation of the worldsheet conformal anomaly. Here, ten is a KO-theoretic invariant, defined only modulo eight, of the real structure of a product spectral triple; it counts no spatial directions, and involves no strings, no extended objects, no worldsheet, and no anomaly to cancel. PST arrives at ten by Connes–Bott periodicity applied to a distinction-first substrate.

In this framework the instantiated geometry is four-dimensional and Lorentzian; the remaining six of the KO-bookkeeping live in a finite noncommutative internal factor with no spatial extent, the algebraic home of the gauge structure and the matter content. PST goes beyond a substrate-level grounding of the Connes–Chamseddine programme [49]: the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ that programme posits is, in PST, *derived* from the primitive distinction structure (§8.8). The substrate's parity-selected real structure (Computation 3) delivers the $\text{Cl}(0,6)$ chirality grading exactly (Computation 19: $\gamma_K = i\omega$ to machine precision in the Jordan–Wigner representation), and Furey 2014 [102] calculates A_F as the commutant of the matter representation on the resulting minimal left ideal. The remainder of this paper develops

the structures that follow (general relativity via the spectral action, the gauge group as the unimodular unitaries of A_F , the matter content of one generation, and quantum mechanics), and isolates the two remaining open pieces in Section 14.2.

1.5. 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 is called this 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, the substrate’s algebraic structure organises itself into two qualitatively different pieces. The Mosco limit of the Boolean Dirichlet form produces a four-dimensional Lorentzian spacetime $M = \mathbb{R} \times S^3$ directly; there is no ‘dimensional reduction’ from a higher-dimensional manifold. Separately, the substrate’s internal algebraic content is the finite noncommutative factor

$$A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C}),$$

an algebraic object with no spatial extent. This algebra identity is derived from P1–P3 alone, via three computationally verified pieces (§8.8): Computation 3 verifies that the substrate’s parity-selected real structure on an odd number of distinction bits delivers $\gamma_K^2 = +I$ with the Connes Lorentzian KO-6 sign triple $(\varepsilon, \varepsilon', \varepsilon'') = (+, +, -)$; Computation 19 verifies that γ_K equals the $\text{Cl}(0,6)$ chirality grading $i\omega$ exactly in the Jordan-Wigner representation, so the chirality projector $f = (I + i\omega)/2$ is delivered directly by the substrate; Furey 2014 [102] then derives A_F as the commutant of the matter representation on the resulting minimal left ideal. The earlier draft of this paper carried the algebra identity as Conjecture 1; that conjecture is now closed, with no appeal to the Chamseddine-Connes-Marcolli classification. The foundational object PST identifies is the product of M and F , a real spectral triple of total KO-dimension ten (Section 1.4); the four, six, and ten are the minimal split compatible with a Lorentzian fermion action, given the verified parity-selected substrate KO-6 (Computation 3) and the verified KO-4 spin structure of the Mosco-limit spacetime (Computation 4).

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 [18, 19] has approached this question from within the instantiated domain; PST proposes the precausal domain as the natural resolution point.

1.6. 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 falsifiable Casimir correction [11, 12, 13, 14, 15] whose d^{-6} scaling is a necessary signature (shared with other single-length-scale vacuum modifications) and whose parameter-free coefficient $\xi = 90/\pi^2$ fixes the diagnostic amplitude ξd_0^2 that selects PST, measurable at submicron plate separations.

1.7. The foundational postulates of PST

PST rests on three independent foundational postulates (P1–P3), plus a fourth formal statement (P4) that is derived from P1–P3 via Mosco convergence of Boolean Dirichlet forms (proved in Section 7.6). The postulates are stated in terms of a *direction algebra* V_7 , a 7-dimensional vector space over \mathbb{R} equipped with the octonion vector product. The dimension is not chosen: it is fixed by Hurwitz’s classification (theorem below).

H (Hurwitz constraint – structural input, not a postulate). The only normed division algebras over \mathbb{R} are $\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}$, of total dimension 1, 2, 4, 8. The corresponding *vector parts* (orthogonal complements of the identity) have dimensions 0, 1, 3, 7. The maximal vector dimension compatible with norm-multiplicativity is therefore 7, realised uniquely by the octonion vector product on V_7 . We adopt this maximal direction algebra as the carrier of directional content in PST.

P1 (The Distinction Primitive, with directional content). There exists a nonspatiotemporal set D of primitive properties. Each property $a \in D$ carries inherent directional content $v(a) \in V_7$. The substrate is the structure (D, δ, v) where: (i) δ is the symmetric irreflexive distinction relation, and (ii) $v: D \rightarrow V_7$ is the direction assignment. No geometric, causal, or temporal structure is presupposed; the directional content in V_7 is the only structural enrichment of bare distinguishability and is forced by H.

P2 (Asymmetric Vector Tension). The tension functional $T: \mathcal{P}(D) \rightarrow V_7$ is V_7 -valued. It is asymmetric under complementation: $T(D \setminus C) \neq -T(C)$ generically. Equivalently, the total substrate tension $T(D)$ does not vanish; it provides a substrate-level directional bias in V_7 . This bias is the source of the preferred direction $\hat{\tau} = T(D)/|T(D)|$ that is selected at modal sublimation. Asymmetric vector tension is the sole source of directedness, including the chiral direction that organises both the colour gauge group and the generation count (§8.17).

P3 (The Landau–Ginzburg Modal Potential on V_7). The order parameter $\psi: \mathcal{C} \rightarrow V_7$ is V_7 -valued. Its dynamics near the modal threshold are governed by the G_2 -invariant Landau–Ginzburg functional

$$\mathcal{F}[\psi, \varepsilon] = \int_{\mathcal{C}} \left[-\varepsilon \langle \psi, \psi \rangle + \frac{1}{2} \langle \psi, \psi \rangle^2 + c |\nabla_{\mathcal{C}} \psi|_{V_7}^2 \right] d\mu, \quad (1)$$

where $G_2 = \text{Aut}(\mathbb{O})$ is the automorphism group of the octonion product (which acts on V_7 in its 7-dim irreducible representation), $\langle \cdot, \cdot \rangle$ is the canonical inner product on V_7 , and $\varepsilon = |T(C)| - \tau$ is the scalar excess tension. The gradient term is the Boolean Dirichlet form on $\{0, 1\}^D$ extended pointwise on V_7 values:

$$|\nabla_{\mathcal{C}} \psi|_{V_7}^2 \equiv \sum_{\substack{C, C' \in \{0, 1\}^D \\ |C \Delta C'| = 1}} \langle \psi(C) - \psi(C'), \psi(C) - \psi(C') \rangle. \quad (2)$$

The potential admits a continuous family of degenerate minima at $|\psi| = r_0(\varepsilon)$ parameterised by the unit direction $\hat{\psi} \in S^6 \subset V_7$. Modal sublimation locks $\hat{\psi}$ to $\hat{\tau} = T(D)/|T(D)|$ – not by spontaneous symmetry breaking, but by direct determination from the substrate’s directional bias (P2).

P4 (The Laplace-Type Projection, derived). The projection kernel $\tilde{K}(x, y; t)$ is, to leading order in t , the heat kernel of a Laplace-type operator $\Pi(L)$ on $L^2(M, \text{dvol}_g)$ whose principal symbol is the emergent metric:

$$\Pi(L) = -g^{\mu\nu} \nabla_\mu \nabla_\nu + (\text{lower order}), \quad (3)$$

so that $\tilde{K}(x, y; t) = \langle x | e^{-t\Pi(L)} | y \rangle$. P4 is derived from P1–P3 by Mosco convergence of the V_7 -valued Boolean Dirichlet form to the standard Laplace–Beltrami form on $M = \mathbb{R} \times S^3$ along the selected direction $\hat{\tau}$ (Section 7.6).

Three independent postulates. H is a mathematical classification theorem, not a postulate. P1–P3 are the three independent postulates of PST: distinguishability with V_7 -directional content, asymmetric vector tension, and the G_2 -invariant Landau–Ginzburg potential. P4 follows from P1–P3 via Mosco convergence.

Recovery of complex-scalar formulation. The complex-scalar Landau–Ginzburg formulation of PST’s earlier exposition is recovered as a 2-dimensional slice of the present V_7 -formulation: choosing any 2-plane $\mathbb{C}_{\hat{\tau}, \hat{\eta}} \subset V_7$ spanned by $\hat{\tau}$ and a perpendicular unit direction $\hat{\eta} \perp \hat{\tau}$ projects the present potential to the familiar complex-scalar form $\mathcal{F}_{\mathbb{C}}[\psi_{\mathbb{C}}, \varepsilon]$ with $\psi_{\mathbb{C}} = \langle \psi, \hat{\tau} \rangle + i \langle \psi, \hat{\eta} \rangle$. The ring-of-minima S^1 vacuum manifold of the scalar formulation is the projection of the S^6 vacuum manifold to that 2-plane. All quantitative predictions involving the scalar-formulation ring are preserved under the projection; the additional 5 directions perpendicular to $\mathbb{C}_{\hat{\tau}, \hat{\eta}}$ carry the gauge-theoretic content (colour SU(3) and the three generations, §8.17).

The Bernoulli measure $\mu = \bigotimes_{a \in D} \text{Bern}(\frac{1}{2})$ is not a postulate; it is the unique measure on $\mathcal{P}(D)$ invariant under relabelling and under complementation (Section 3). The Lorentzian signature $(-, +, +, +)$ is derived from hyperbolic well-posedness of the Goldstone wave equation along $\hat{\tau}$ (Computation 2). See Section 14.2 for the full accounting and for the open research that remains.

1.8. 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 8 derives the Standard Model gauge group and establishes gauge anomaly cancellation as a structural theorem. Section 9

develops the physical predictions of the theory, including a quantitative Casimir correction. Section 10 grounds Newton’s constant as the gravitational term of the spectral action and identifies the discreteness scale d_0 with the Landau-Ginzburg coherence length of the modal vacuum, establishing the Casimir correction (the d^{-6} scaling is a necessary signature; the parameter-free coefficient $\xi = 90/\pi^2$ fixes the diagnostic amplitude ξd_0^2 that selects PST) as a concretely detectable experimental signal. Section 11 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 12 positions PST in relation to existing frameworks. Section 13 concludes. Section 14 summarises the structural and mathematical problems the theory resolves and the open research that remains. Appendix 14.3 is a consolidated table of the symbols recurring through the paper, grouped by sector, each with the section where it is first introduced.

1.9. Layman’s Summary

Physics has produced 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

2.1. The case for ontological minimality

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.

2.2. 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 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.

2.3. Mathematical formulation

The substrate is the triple (D, δ, v) :

- D is a set of primitive properties.
- $\delta: D \times D \rightarrow \{0, 1\}$ is the distinction relation, symmetric and irreflexive:

$$\delta(a, b) \implies \delta(b, a) \tag{4}$$

$$\neg \delta(a, a) \tag{5}$$

- $v: D \rightarrow V_7$ is the **direction assignment**: each property $a \in D$ has inherent directional content $v(a) \in V_7$ in the 7-dimensional direction algebra.

No further axioms are imposed on δ . There is no transitivity requirement, no totality requirement, no metric on D , no topology, and no ordering. The direction assignment v replaces a labelled-set structure with a V_7 -vector content; without it, distinctions would carry magnitude only (“this property is distinct”) with no record of *what kind* of distinction. With v , each distinction is a vector in the maximal Hurwitz direction algebra, structured by the octonion vector product.

The 7-dimensional direction algebra V_7 . $V_7 \cong \mathbb{R}^7$ as a vector space, equipped with:

- a positive-definite inner product $\langle \cdot, \cdot \rangle$,
- a bilinear vector product $\times: V_7 \times V_7 \rightarrow V_7$ satisfying $\langle u \times v, u \rangle = \langle u \times v, v \rangle = 0$ and $|u \times v|^2 = |u|^2|v|^2 - \langle u, v \rangle^2$.

This is the unique 7-dimensional algebra satisfying these properties (Hurwitz); equivalently, it is the vector part of the octonions \mathbb{O} under octonion multiplication. The automorphism group of V_7 (preserving both the inner product and the vector product) is the exceptional Lie group G_2 , of real dimension 14. The choice of dimension 7 for the direction algebra is forced: no other dimension supports a maximal normed-product structure over \mathbb{R} .

Configurations and the direction content of a configuration. A **configuration** $C \subseteq D$ is a collection of properties considered jointly. The space of configurations is $\mathcal{C} = \mathcal{P}(D)$. The total direction content of a configuration is the sum

$$v(C) = \sum_{a \in C} v(a) \in V_7. \tag{6}$$

This is the natural V_7 -valued aggregate of the directional content of the properties in C . The asymmetric tension functional of P2 is structurally related to (but distinct from) $v(C)$: $T(C)$ encodes the *magnitude and directional bias* of the configuration’s tension, with the magnitude governing modal sublimation and the direction inheriting from v via the asymmetric tension

construction of Section 2.4.

It is within the configuration space \mathcal{C} that the modal structure of PST is expressed. These are not temporal dynamics; they concern what is possible, what is necessary, and what cannot remain uninstantiated.

2.4. 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 pre-causal 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 pre-causal 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 V_7 -valued map on configurations:

$$T: \mathcal{C} \rightarrow V_7. \quad (7)$$

The V_7 -vector $T(C)$ encodes both the *magnitude* of the configuration's modal imbalance, $|T(C)| = \sqrt{\langle T(C), T(C) \rangle}$, and its *directional content* in the 7-dimensional direction algebra. The magnitude is strictly positive on nontrivial configurations (those with at least one pair of distinct properties); the direction inherits from the directional content $v: D \rightarrow V_7$ assigned to the elementary properties (P1). A configuration with no distinct properties (the null

configuration) carries $T = 0$ trivially.

Representing tension as a V_7 -valued map carries strictly more structure than a real-scalar representation. The magnitude provides a partial ordering on configurations (any two configurations can be compared by $|T|$), and this partial ordering is what the variational analysis of Section 3 uses. The directional content provides the substrate-level information that propagates through the modal-sublimation chain to fix the colour gauge group and the generation count (§8.17). The V_7 -valued representation is forced by the Hurwitz constraint H and the directional content of P1: a real-scalar representation would discard the directional information $v(a)$ encoded in the elementary properties of D .

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.

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

$$T(\bar{C}) \neq -T(C) \tag{8}$$

generically. Equivalently, the total substrate tension

$$T(D) = T(C) + T(\bar{C}) \neq 0 \tag{9}$$

does not vanish for generic substrates. This is the vector generalisation of the scalar asymmetry condition $T(C) \neq T(\bar{C})$ of earlier exposition: in the V_7 -valued formulation, $T(D)$ is the substrate's net directional bias in V_7 , a V_7 -vector that does not cancel between a configuration and its complement.

If $T(D) = 0$ (the symmetric case), then $T(\bar{C}) = -T(C)$ for every configuration, and the substrate has no preferred direction in V_7 . PST asserts the generic case $T(D) \neq 0$: the substrate has a structural net direction, and this net direction $\hat{\tau} = T(D)/|T(D)|$ is what modal sublimation locks into (P3). The net imbalance is therefore always a nonzero V_7 -vector:

$$\Delta T(C) := T(C) - T(\bar{C}) \neq 0 \quad \text{in } V_7. \tag{10}$$

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 V_7 -valued character of T encodes that the imbalance has both magnitude and directional content; the magnitude is what eventually drives modal sublimation, and the direction is what fixes the chiral structure of the instantiated geometry.

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. 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* [59] 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* [65] 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 pre-causal 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.

2.5. Layman’s Summary

This section establishes the first and most minimal starting point of the theory: the fact that things can be different from each other. No space, no time, no energy, no forces, and no laws of nature are assumed here. 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. The Landau-Ginzburg form of the modal potential and the Laplace-type property of the projection kernel are introduced as explicit postulates in later sections (see Section 1.7).

3. The Modal Threshold

3.1. Why a modal threshold must exist

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 pre-causal 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, the magnitude $|T(C)| = \sqrt{\langle T(C), T(C) \rangle}$ of its V_7 -valued asymmetric tension increases. This monotonicity is not derivable from the axioms of (D, δ, v) alone; it is a structural constraint on T that PST adopts as part of its encoding. Three structural constraints on T are required:

$$|T(C)| \geq 0, \quad |T(C)| = 0 \Leftrightarrow C \text{ has no distinct pairs} \quad (11)$$

$$|T(C \cup \{a\})| \geq |T(C)| \quad \forall a \notin C \quad (12)$$

$$T(\bar{C}) \neq -T(C) \quad \text{generically} \quad (13)$$

Equation (11) is magnitude positivity; equation (12) is magnitude monotonicity; equation (13) is complement vector asymmetry, the constraint that the total substrate tension $T(D) = T(C) + T(\bar{C})$ does not vanish, providing the substrate-level directional bias \hat{r} that P2 names. A minimal concrete example satisfying all three simultaneously is the **induced edge count**:

$$T_{\text{ex}}(C) = |\{(a, b) : a, b \in C, a \neq b, \delta(a, b)\}| \quad (14)$$

Equation (11) holds for any configuration with at least two distinct properties. Equation (12) holds by inspection: adding property a to C increases the count by $|\{b \in C : \delta(a, b)\}|$. Equation (13) holds whenever the induced subgraphs on C and \bar{C} have different edge counts, which is generic for non-regular distinction structures. Equations (11)–(13) 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.

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.

3.2. Boundary condition

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

$$T(C) \geq \tau \implies C \text{ cannot remain uninstantiated} \quad (15)$$

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**:

$$\varepsilon(C) = T(C) - \tau \quad (16)$$

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.

3.3. 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 $\psi(C) \in V_7$ as the **modal order parameter**: a V_7 -valued field over configuration space encoding the modal status of each configuration. The value $\psi = 0$ denotes the uninstantiated phase. Nonzero values of the magnitude $|\psi| = \sqrt{\langle \psi, \psi \rangle}$ indicate degrees of instantiation, with larger $|\psi|$ corresponding to deeper settlement into the instantiated vacuum. The unit direction $\hat{\psi} = \psi/|\psi|$ on $S^6 \subset V_7$ carries the directional content of the instantiation event. The V_7 -valued representation inherits from the directional content of P1: an order parameter for asymmetric vector tension must take values in the same direction algebra.

The full G_2 symmetry of V_7 (rotations preserving the inner product and the vector product) is the symmetry under which the modal potential is invariant; the \mathbb{Z}_2 symmetry $\psi \rightarrow -\psi$ of earlier exposition is the antipodal involution on S^6 , a G_2 -equivariant special case. Each of the S^6 directions is a candidate “which way” the vacuum will settle; modal sublimation locks $\hat{\psi}$ to the substrate’s net direction $\hat{\tau} = T(D)/|T(D)|$ (P2).

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} , taking values in V_7 : it assigns to each precausal configuration a V_7 -vector 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 magnitude τ and the substrate-determined direction $\hat{\tau}$ from the structure of the modal potential.

3.4. 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 \S carries no preferred orientation in V_7 at the pre-sublimation level: all unit directions on $S^6 \subset V_7$ are candidate vacuum positions, and the modal potential must be invariant under the action of the automorphism group G_2 of V_7 . This G_2 -symmetry is the natural generalisation of the discrete $\psi \rightarrow -\psi$ reflection symmetry of a real-scalar order parameter. The functional must therefore be invariant under all G_2 -rotations of $\psi \in V_7$. This eliminates all G_2 -non-invariant tensor expressions; the only admissible polynomials in ψ are functions of the G_2 -invariant inner product $\langle \psi, \psi \rangle$, and the gradient G_2 -invariant $|\nabla_{\mathcal{C}}\psi|_{V_7}^2$.

Taylor expansion near the threshold. Close to the threshold, $|\psi|$ is small and \mathcal{F} can be expanded as a power series in $\langle \psi, \psi \rangle$. The leading term allowed by G_2 -symmetry is $\langle \psi, \psi \rangle$, 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 magnitude, which vanishes precisely at the threshold.

The next allowed term is $\langle \psi, \psi \rangle^2$ 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-in-magnitude 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 ($\langle \psi, \psi \rangle^3, \dots$) 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 $\int [-\varepsilon \langle \psi, \psi \rangle + b \langle \psi, \psi \rangle^2] d\mu$ 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 $|\psi(C) - \psi(C')|_{V_7} \rightarrow 0$ as $|C \Delta C'| \rightarrow 0$, where the norm is taken on V_7 . 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 V_7 -squared gradient $|\nabla_{\mathcal{C}}\psi|_{V_7}^2 = \sum_{C \sim C'} \langle \psi(C) - \psi(C'), \psi(C) - \psi(C') \rangle$ 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 [58] 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 $d_{\Delta}(C, C') = |C\Delta C'|$, where $C\Delta C'$ 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 $\text{Aut}(D, \delta)$), 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 provides is, in fact, all three requirements — (a), (b), and (c) — once the Bernoulli measure μ (derived in the next subsection) is in hand. The construction is as follows.

The Boolean gradient. For each locus $a \in D$, define the **Boolean derivative** of ψ at a by

$$(\partial_a\psi)(C) := \psi(C \cup \{a\}) - \psi(C \setminus \{a\}), \quad (17)$$

where $C \cup \{a\}$ is the configuration C with locus a activated, and $C \setminus \{a\}$ is C with a inactivated. This is the unique finite-difference operator compatible with the graph structure of $(\mathcal{C}, d_{\Delta})$: two configurations are adjacent in $(\mathcal{C}, d_{\Delta})$ iff they differ by exactly one locus activation, and $\partial_a\psi$ measures the variation of ψ across the unique edge incident on both $C \cup \{a\}$ and $C \setminus \{a\}$.

The full **Boolean gradient** is the vector of all these derivatives:

$$|\nabla_C \psi|^2(C) := \frac{1}{4} \sum_{a \in D} |\psi(C \cup \{a\}) - \psi(C \setminus \{a\})|^2. \quad (18)$$

L² differentiable structure. The Bernoulli measure μ on $\mathcal{P}(D)$ turns $L^2(\mathcal{P}(D), \mu)$ into a Hilbert space. For any $\psi \in L^2(\mathcal{P}(D), \mu)$, the quantity $\mathbb{E}_\mu[|\nabla_C \psi|^2]$ is finite, because each $\partial_a \psi$ is bounded in L^2 whenever ψ is. The gradient operator $\nabla_C : L^2(\mathcal{P}(D), \mu) \rightarrow \bigoplus_{a \in D} L^2(\mathcal{P}(D), \mu)$ is a bounded linear operator, and the Sobolev space $W^{1,2}(\mathcal{P}(D), \mu) = \{\psi : \mathbb{E}_\mu[|\nabla_C \psi|^2] < \infty\}$ is a Hilbert space with inner product $\langle \psi, \phi \rangle_{1,2} = \langle \psi, \phi \rangle_{L^2} + \sum_a \langle \partial_a \psi, \partial_a \phi \rangle_{L^2}$ [54]. The inner product on tangent directions (the role of item (c)) is the $L^2(\mu)$ -inner product on the Boolean derivatives, which requires no additional primitive beyond μ and d_Δ .

Poincaré inequality. The Efron–Stein inequality [53] states: for any $\psi \in W^{1,2}(\mathcal{P}(D), \mu)$,

$$\text{Var}_\mu(\psi) \leq \frac{1}{4} \sum_{a \in D} \mathbb{E}_\mu[|\partial_a \psi|^2] = \mathbb{E}_\mu[|\nabla_C \psi|^2], \quad (19)$$

with equality when $\mu = \bigotimes_a \text{Bern}(1/2)$ and ψ is linear in the indicator functions $\mathbf{1}_a$ [54]. This Poincaré inequality guarantees that the gradient term in \mathcal{F} controls the variance of ψ about its mean: configurations far apart in d_Δ are forced to have similar ψ values by the cost $c |\nabla_C \psi|^2$, which is exactly the coherence requirement the gradient term was introduced to enforce.

Resolution of the apparent gap. The differentiable structure (b) is provided by the $L^2(\mu)$ Hilbert space structure, whose existence follows from the Bernoulli measure μ derived in the next subsection. The inner product on tangent directions (c) is the $L^2(\mu)$ -inner product on Boolean derivatives, derived from d_Δ alone. No additional topological or geometric primitive is required. The gradient structure problem is thereby resolved within the existing primitive structure of PST.

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** $\text{Aut}(D, \delta)$, consisting of all bijections $f : D \rightarrow D$ 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 $\text{Aut}(D, \delta)$ 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:

$$\mu(C) = 2^{-n} \quad \text{for all } C \in \mathcal{C} \quad (20)$$

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 :

$$\mu(\{C \in \mathcal{C} : C_0 \subseteq C\}) = 2^{-|C_0|} \quad \text{for all finite } C_0 \subseteq D \quad (21)$$

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**:

$$\mathcal{F}(\psi, C) = \int_{\mathcal{C}} [a(T(C)) \psi^2 + b \psi^4 + c |\nabla_{\mathcal{C}} \psi|^2] d\mu(C) \quad (22)$$

where $\nabla_{\mathcal{C}}$ denotes the gradient over configuration space (not spacetime), μ is the canonical Bernoulli measure of equation (21), 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:

$$\mathcal{F}(\psi) = -\varepsilon \psi^2 + \frac{1}{4} \psi^4, \quad \varepsilon = T(C) - \tau \quad (23)$$

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 (22) 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, $c|\nabla_C\psi|^2$ 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 treated 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.

3.5. 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 $-\varepsilon\langle\psi, \psi\rangle$ 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 $b\langle\psi, \psi\rangle^2$ is the nonlinear saturation term: it is always positive and ensures the functional remains bounded below. The gradient term $c|\nabla_C\psi|_{V_7}^2$ penalises rapid variation of ψ across configuration space. Distributing $\delta/\delta\psi$ across these three terms and collecting, the modal equation of state is the V_7 -vector equation:

$$-2\varepsilon\psi + 4b\langle\psi, \psi\rangle\psi - 2c\nabla_C^2\psi = 0 \quad (24)$$

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

Below the threshold, $|T(C)| < \tau$: $\varepsilon < 0$. The only solution to the uniform equation is $\psi = 0 \in V_7$. Since $\delta^2\mathcal{F}/\delta\psi^2|_{\psi=0} = -2\varepsilon\mathcal{K}_{V_7} > 0$ (positive multiple of the identity), this is a local minimum in every direction of V_7 . The configuration remains uninstantiated. This is the *precausal phase*.

Above the threshold, $|T(C)| > \tau$: $\varepsilon > 0$. The solution $\psi = 0$ is now a local maximum of \mathcal{F} in every V_7 -direction. The quadratic term falls with increasing $|\psi|$; the quartic term turns \mathcal{F} upward again at large $|\psi|$. The result is a G_2 -invariant Mexican-hat profile: a continuous family of degenerate minima at fixed radius, parameterised by the unit direction $\hat{\psi} \in S^6 \subset V_7$.

The location of these minima requires explicit computation. Restricting to spatially uniform configurations, the equation of state reduces to $-2\varepsilon\psi + 4b|\psi|^2\psi = 0$, i.e. $2\psi(-\varepsilon + 2b|\psi|^2) = 0$. The root $\psi = 0$ is the maximum. The nontrivial roots satisfy $|\psi|^2 = \varepsilon/(2b)$, giving the vacuum 6-sphere

$$\mathcal{V} = \{ \psi \in V_7 : |\psi| = r_0 \} \cong S^6, \quad r_0 = \sqrt{\frac{|T(C)| - \tau}{2b}} = \sqrt{\frac{\varepsilon(C)}{2b}}. \quad (25)$$

Modal sublimation locks $\hat{\psi}$ to the substrate-determined direction $\hat{\tau} = T(D)/|T(D)|$ on \mathcal{V} (P2). The earlier exposition’s two “symmetric minima at $\psi = \pm r_0$ ” is the antipodal pair $\pm r_0\hat{\tau}$ in the 1-dim projection along $\hat{\tau}$; the full vacuum manifold has 5 additional directions on S^6 orthogonal to $\hat{\tau}$ that carry the gauge-theoretic content (§8.17). 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 $\psi \rightarrow -\psi$ 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.

3.6. The bifurcation point and the derivation of τ

The critical tension magnitude τ is the value at which the qualitative character of \mathcal{F} changes: for $|T(C)| < \tau$, the uninstantiated phase $\psi = 0 \in V_7$ is the unique minimum; for $|T(C)| > \tau$, it is a maximum and a 6-sphere of new minima appears. The moment of change is the **bifurcation point**, the magnitude of T at which $\psi = 0$ transitions from stable to unstable.

The characterisation of critical points as minima or maxima requires the second variation. For a V_7 -valued field, the second functional derivative is a quadratic form on V_7 (a 7×7 symmetric tensor). If the form is positive-definite, the functional is locally convex at ψ_0 in every V_7 -direction; the configuration is at a stable minimum. If indefinite or negative-definite, at least one direction in V_7 is unstable.

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 $\delta\mathcal{F}/\delta\psi = -2\varepsilon\psi + 4b|\psi|^2\psi$ once more with respect to ψ and evaluating at $\psi = 0$ eliminates the cubic-in- ψ terms entirely. What remains is purely the contribution of the quadratic term, proportional to the identity on V_7 :

$$\left. \frac{\delta^2\mathcal{F}}{\delta\psi^2} \right|_{\psi=0} = -2\varepsilon(|T|)\mathbb{K}_{V_7} = 2(\tau - |T|)\mathbb{K}_{V_7} \quad (26)$$

This is positive-definite for $|T| < \tau$ (i.e., $\varepsilon < 0$), zero for $|T| = \tau$, and negative-definite for $|T| > \tau$ (i.e., $\varepsilon > 0$). The G_2 -isotropy of the second variation at $\psi = 0$ is direct: the form is a scalar multiple of the identity, hence invariant under all V_7 -rotations. 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 $\delta^2\mathcal{F}/\delta\psi^2|_{\psi=0} = 0$, 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 magnitude below which $\psi = 0$ is a stable minimum and above which it is not. This is precisely the infimum of the set of tension magnitudes for which the second variation vanishes:

$$\tau = \inf \left\{ |T(C)| : \left. \frac{\delta^2 \mathcal{F}}{\delta \psi^2} \right|_{\psi=0} = 0 \right\} \quad (27)$$

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 $\varphi = \pm\sqrt{2\varepsilon}$; the configuration at $\varphi = 0$ is now a local maximum and cannot remain there. The dashed curve is the reference potential at $\varepsilon = 0$. The four key relations carried by this figure are: (1) the Landau-Ginzburg modal potential $F(\psi, \varepsilon) = -\varepsilon \langle \psi, \psi \rangle + \frac{1}{4} \langle \psi, \psi \rangle^2$ with quartic coefficient $b = \frac{1}{4}$ fixed by the LG form; (2) excess tension $\varepsilon = |T(\mathcal{C})| - \tau$; (3) the vacuum manifold radius $|\psi^*| = \sqrt{2\varepsilon}$ on V_7 with topology $\mathcal{V} \cong S^6$ (a 6-sphere of degenerate directions in V_7 , projecting onto S^1 in any 2-plane slice of V_7); and (4) the projection map $\Phi : \mathcal{S} \rightarrow \mathcal{G}$ from the precausal substrate to instantiated geometry.

3.7. 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

4.1. The character of the transition

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.

4.2. 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.

4.3. 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.

4.4. 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.

4.5. 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.

4.6. 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 \mathbf{S}** has as objects all configurations $C \in \mathcal{C}$ with $T(C) \geq \tau$, and as morphisms distinction preserving maps $f : C \rightarrow C'$. The **geometric category \mathbf{G}** has as objects Lorentzian manifolds (M, g) and as morphisms smooth maps preserving the causal structure.

Modal sublimation is the functor:

$$\Phi : \mathbf{S} \rightarrow \mathbf{G} \tag{28}$$

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

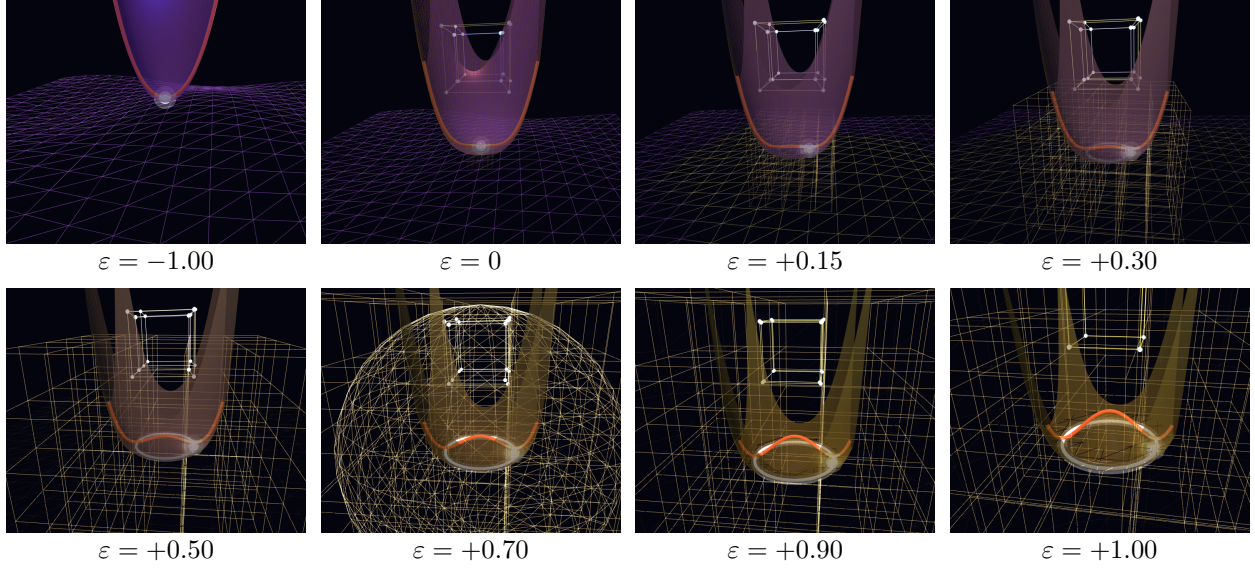


Figure 2: Visual model of modal sublimation across eight values of the excess tension $\varepsilon = T(C) - \tau$ (Section 3): $\varepsilon < 0$ precausal, $\varepsilon \geq 0$ instantiated. Sampling is dense on the instantiated side, where the geometric content emerges; a single precausal reference frame at $\varepsilon = -1.00$ shows the pre-threshold state from which everything else unfolds. 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-left frame* ($\varepsilon = -1.00$): 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 a flat proto-manifold grid is shown for reference. *Threshold and early instantiation* ($\varepsilon = 0$ to $\varepsilon \approx +0.30$): the bowl flattens and its centre rises, the high-dimensional cube descends toward the bowl, the spacetime lattice crystallises, and the ring of minima begins to form at the base. *Mid to deep instantiation* ($\varepsilon \approx +0.50$ to $+1.00$): the potential has acquired a continuous 6-sphere of degenerate minima (the vacuum manifold $\mathcal{V} \cong S^6 \subset V_7$); the figure displays the 2-plane projection visible as an orange ring $\mathcal{V}_{\hat{\eta}} \cong S^1$ in any $\hat{\tau}$ -perpendicular slice; the spatial slice closes, shown by the S^3 dome enveloping the local chart; and the finite internal factor $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ consolidates as the cube snaps onto the spacetime grid.

4.7. Explicit construction of Φ

The projection functor Φ inherits the post-sublimation directional structure of the substrate: it is constructed *relative to* the $\hat{\tau}$ -direction that modal sublimation locks (P2). The construction below produces a spacetime M and a realisation ρ together with an internal-factor identification along $\hat{\tau}$; the same construction at a different $\hat{\tau}' \in S^6$ would produce a G_2 -rotated copy of the same Φ , with the rotation an unphysical relabelling. Once $\hat{\tau}$ is fixed by the substrate's directional bias, Φ is determined up to the gauge freedom described below.

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 $\rho : D \rightarrow M$ 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 $\rho : D \rightarrow X$, 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_ρ 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_ρ 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** $\rho : D \rightarrow M$ 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. The metric is then:

$$g_{\mu\nu}(\rho(a), \rho(b)) = \eta_{\mu\nu} \cdot \text{sgn}(T(\{a, b\}) - \tau_{\text{null}}) \cdot |T(\{a, b\}) - \tau_{\text{null}}|^{1/2} \quad (29)$$

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 forced as a theorem by hyperbolic well-posedness of the Goldstone wave equation (Computation 2, Section 14.2): the Cauchy problem for $\square_g \theta = 0$ admits a well-posed initial-value formulation and a conserved Noether current only for signature $(1, n-1)$. The bifurcation of tension values around τ_{null} remains the algebraic mechanism inside the LG variational structure; what Computation 2 shows is that the sign cut required for that bifurcation is not an independent axiom but a consequence of the well-posedness condition.

Two clarifications are necessary. First, equation (29) 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 (29) along the chord, analogous to recovering a Riemannian metric as the Hessian of a distance function. Second, the magnitude $|T(\{a, b\}) - \tau_{\text{null}}|^{1/2}$ 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 (134) blows up precisely at the modal null threshold rather than at $T = 0$.

Step 3, Action on morphisms. For $f : C \rightarrow C'$ in \mathbf{S} , $\Phi(f) : (M, g) \rightarrow (M', g')$ 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 :

$$\Pi(T(C))(x) = \int_{\mathcal{C}} T(C') \cdot K(x, C') d\mu(C') \quad (30)$$

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 [48]: 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:

$$T_{\mu\nu}^{(\text{iso})}(x) = \Pi(T(C))(x) \cdot g_{\mu\nu}(x) \quad (31)$$

This isotropic scalar form $T_{\mu\nu}^{(\text{iso})} = \Pi(T(C)) g_{\mu\nu}$ 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. 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.

4.8. Uniqueness of Φ and the origin of diffeomorphism invariance

The construction proceeds via an explicit choice of realization map $\rho : D \rightarrow M$, 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 (29): there are no further choices once the realization map is fixed. Now suppose $\rho' = f \circ \rho$ for some smooth bijection $f : M \rightarrow M'$. 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 (29) 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 natural question of 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 note on the metric signature. The argument in Step 2 assigned the Minkowski signature $(-, +, +, +)$ through equation (29), by bifurcating pairwise tensions about a null threshold τ_{null} that separates spacelike from timelike separations. This determines the signature *given* τ_{null} : once a sign cut is supplied, $T > \tau_{\text{null}}$ is spacelike and $T < \tau_{\text{null}}$ is timelike, by construction. It does *not* determine that exactly one direction lies below τ_{null} , spectral asymmetry of the tension spectrum about a threshold (the standard form of which is $T(C) \neq T(\bar{C})$, complement-asymmetry) is consistent with zero, one, or several eigenvalues lying below any given cut, so it does not by itself fix the count of timelike directions. Equivalently: in the multidimensional-scaling reframing of the realisation map (Section 14.2, item (1)), any S_D -invariant Gram construction over real features is positive-semidefinite by elementary linear algebra, so the negative eigenvalue required for one timelike direction cannot arise from S_D -symmetry alone; it must be supplied by an indefinite bilinear form imposed by hand. The Lorentzian signature $(-, +, +, +)$ is fixed as a *theorem* via Computation 2 of Section 14.2: hyperbolic well-posedness of the Goldstone wave equation forces $(p, q) = (1, n - 1)$, and the $U(1)$ Noether-current conservation that produces the angular-momentum identification requires this same Lorentzian signature. The macroscopic dimension $n = 4$ that combines with the signature to give the four-dimensional Lorentzian spacetime $M = \mathbb{R} \times S^3$ follows from KO additivity given the verified spectral-triple structure (§1.4): the substrate’s parity-selected KO-6 (Computation 3) and the product’s KO-2 mod 8 (Connes Lorentzian value) together force M at KO-dimension four. PST is then the theory that derives everything downstream from $(D, \delta) + \text{LG} +$ the realisation map ρ , with both signature and macroscopic dimension outputs rather than inputs.

4.9. 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

5.1. Causality as coinstantiated structure

The previous section established that modal sublimation is a nontemporal, nonmechanistic transition: the functor $\Phi : \mathbf{S} \rightarrow \mathbf{G}$ 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 [34]. In causal set theory, the discrete partial order is the primary structure from which the continuum metric is to be recovered [4, 35]. 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 [36]. 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.

5.2. 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

structure of g :

$$p \leq q \iff q \in J^+(p, g) \quad (32)$$

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

$$J^+(p, g) = \{ q \in M \mid \exists \text{ future-directed causal curve } \gamma : [0, 1] \rightarrow M, \gamma(0) = p, \gamma(1) = q \} \quad (33)$$

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 (32) has no pre-image in the pre-causal category \mathbf{S} . Formally, there is no morphism or relation in \mathbf{S} that maps to \leq under Φ , because the objects of \mathbf{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 (29). Therefore the following holds:

$$\text{causality} = f(g) = f(\Phi(C)) \quad (34)$$

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

5.3. 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** [37, 38].

In PST this asymmetry is not imposed as an additional datum. It is inherited from the asymmetry of the pre-causal 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 pre-causal 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 [37], 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 $\nabla_C T$ 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.

5.4. 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 is the V_7 -vacuum sphere $\mathcal{V} \cong S^6$. For the $U(1)$ Goldstone-phase and angular-momentum analysis below, the relevant substructure is the 2-plane slice fixed by $\hat{\tau}$ and any perpendicular unit direction $\hat{\eta} \perp \hat{\tau}$: the projected ring $\mathcal{V}_{\hat{\eta}} \cong S^1$. Henceforth in this paper, statements involving “ $\mathcal{V} \cong S^1$ ” refer to this projection, the $U(1)$ -relevant slice of the full S^6 vacuum manifold (P3); the perpendicular five additional directions on S^6 carry the gauge-theoretic content (§8.17). Motion along $\mathcal{V}_{\hat{\eta}}$ is the only stable state the projected vacuum admits in the slice. 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 the structural argument holds that orbital motion at all scales does not require an initial condition that specifies it separately: it is a candidate for the ground state of the causal structure. The explicit Π -mapping connecting this Noether charge to the observed orbital angular momentum of massive bodies is computed in Section 7.18.

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

With coinstantiation established, PST is 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 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 [39, 40]. 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 pre-causal 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?” [41], 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.

5.6. 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** [36] 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 [42] 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 this 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.

5.7. 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

6.1. Overview and the eight-step summary table

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	$T : \mathcal{C} \rightarrow \mathbb{R}^+$, $T(C) \neq T(\bar{C})$ for all C
3	Modal threshold	Bifurcation point τ of $\mathcal{F}(\psi, C)$, eq. (27)
4	Modal sublimation	Functor $\Phi : \mathbf{S} \rightarrow \mathbf{G}$, constructed via ρ and eq. (29)
5	Geometry as spacetime	Lorentzian manifold (M, g) with signature supplied by 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 content with structural spin-statistics from the CAR algebra on the Boolean distinction primitive (Section 8)
8	Angular momentum	Noether charge $L = 2r_0^2 \dot{\theta}$ of the U(1) symmetry of the 2-plane projection $\mathcal{V}_{\hat{\eta}} \cong S^1$ of the V_7 -vacuum sphere $\mathcal{V} \cong S^6$

6.2. 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.

6.3. 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 pre-causal 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.

6.4. 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 (27).

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 .

6.5. 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 $\Phi : \mathbf{S} \rightarrow \mathbf{G}$, 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 (29). 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.

6.6. 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 (29). 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.

6.7. 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 (35). 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 has two parts. Stress-energy and the bosonic field sector are derived in Section 7 from the order parameter ψ . The fermionic sector and the spin-statistics connection are derived separately in Section 8 from the CAR algebra on the substrate’s Boolean distinction primitive: the fermionic excitations inherit the correct anticommutation relations as a structural consequence of the Boolean $\{0, 1\}$ structure, not as a separate postulate. The chain’s claim at this step is therefore that the full matter content (stress-energy, bosonic fields, and fermionic content with structural spin-statistics) is derivable.

6.8. 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.

The order parameter ψ is V_7 -valued (P3), with the full vacuum manifold $\mathcal{V} \cong S^6$. After modal sublimation locks $\hat{\psi}$ to the substrate-determined direction $\hat{\tau} = T(D)/|T(D)|$, the residual fluctuations of ψ orthogonal to $\hat{\tau}$ are tangent to S^6 . In the 2-plane slice fixed by $\hat{\tau}$ and any unit direction $\hat{\eta} \perp \hat{\tau}$, the residual vacuum manifold is the circle $\mathcal{V}_{\hat{\eta}} \cong S^1$, a one-dimensional ring of degenerate minima parameterised by the angle around $\hat{\tau}$. The Landau–Ginzburg functional restricted to this 2-plane is invariant under global phase rotations $\psi \mapsto e^{i\alpha}\psi$ in the slice: 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.

6.9. 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 initial conditions are imposed except as consequences of the vacuum topology. The chain rests on three independent postulates plus one derived structural statement (Section 1.7); the steps between them are structural derivations. 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.

6.10. 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. The chain rests on three independent postulates (P1–P3: the distinction primitive, asymmetric tension, and the Landau-Ginzburg modal potential) together with the Laplace-type projection (P4), which is derived from P1–P3 via Mosco convergence (Section 7.6). The steps between those postulates are structural derivations: given each postulate, the next element of the chain follows necessarily. What is remarkable is not the length of the chain but the precision of its logical structure, the entire background of modern physics traced back to three minimal commitments, with the fourth following as a structural consequence.

7. Energy, Matter, and Fields

7.1. The source problem in General Relativity

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; 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.

7.2. The stress-energy tensor as projected tension

In the isotropic, homogeneous approximation the precise statement is:

$$T_{\mu\nu}^{(\text{iso})}(x) = \Pi(T(C))(x) \cdot g_{\mu\nu}(x) \quad (35)$$

where $\Pi(T(C))(x)$ is the projection of the precausal tension functional onto the point $x \in M$, as defined in equation (30). 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 .

7.3. The field equations as a conservation statement

Einstein's field equations then read:

$$G_{\mu\nu} = 8\pi G \cdot \Pi(T(C)) \quad (36)$$

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 [19] and the entropic gravity programme of Verlinde [18]. 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.

7.4. 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 $c|\nabla_C\psi|^2$ 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: $t \rightarrow -i\tau$ 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 .

7.5. The projected action and the emergence of the Einstein-Hilbert term

Newton’s constant from the spectral action. Newton’s constant is the gravitational coefficient of the spectral action $\text{Tr } f(D/\Lambda)$ [49]. Its heat-kernel asymptotics contribute a cosmological term at order Λ^4 , the Einstein-Hilbert term $\frac{1}{16\pi G} \int R\sqrt{-g}$ at order Λ^2 , and the Yang–Mills and Higgs terms at order Λ^0 . The Seeley–DeWitt heat-kernel computation carried out in this subsection and the next is exactly that Λ^2 coefficient: $1/G \sim f_2 \Lambda^2 \mathcal{N}$, with \mathcal{N} the trace over the finite internal factor and the cutoff $\Lambda \sim 1/d_0$ set by the Landau-Ginzburg coherence scale. The scales r_0 and d_0 are the Landau-Ginzburg vacuum scales of the order parameter, and G emerges as a derived consistency relation among the observables (m_h, M_P, v) rather than a first-principles output.

The field equations (36) 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 $Rg_{\mu\nu}$, 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 (63) contains three terms. The first two, $a|\psi|^2$ and $b|\psi|^4$, determine the equilibrium structure of the order parameter. The third, $c|\nabla_C\psi|^2$, 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 $\nabla_C\psi$ decomposes into two pieces: the Goldstone mode $r_0 \nabla_C\theta$ (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 (the functional form (that the leading curvature term is $\int R\sqrt{-g}$ and not a higher-curvature combination) is established in Section 7.6 via the Seeley–DeWitt expansion; the kernel normalisation coefficient \mathcal{N} is determined explicitly from the Bernoulli-measure projection integral in Section 7.8):

$$\Pi(c|\nabla_{\mathcal{C}}\psi|^2)|_{\psi=r_0} \longrightarrow \frac{r_0^2 c}{16\pi G} \int_M R\sqrt{-g} d^4x + \mathcal{O}(R_{\mu\nu}R^{\mu\nu}, R^2) \quad (37)$$

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 (59). 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 $\gg d_0$. The effective projected action at low energies is therefore:

$$S[g] = \frac{1}{16\pi G} \int_M R\sqrt{-g} d^4x + S_{\text{matter}}[g, \text{fields}] \quad (38)$$

This is the **Einstein-Hilbert action**. The step from the gradient term in \mathcal{F} to equation (38) is structural. The functional form of the projection (that the leading curvature term is $\int R\sqrt{-g} d^4x$ and not a higher-curvature combination) is derived in Section 7.6 via the Seeley–DeWitt heat-kernel expansion (the projection-coefficient question is thereby resolved: the structure of the R term is established here, and its numerical coefficient \mathcal{N} is determined from the Bernoulli-measure projection integral in Section 7.8). The downstream Noether–Lovelock argument is standard and unambiguous given equation (38). Varying equation (38) with respect to $g^{\mu\nu}$ yields, by a standard calculation:

$$\frac{\delta S}{\delta g^{\mu\nu}} = 0 \quad \implies \quad G_{\mu\nu} = 8\pi G T_{\mu\nu} \quad (39)$$

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 (36) is thereby derived, not asserted.

7.6. The Seeley–DeWitt argument for the functional form of equation (37)

The previous subsection deferred the determination of the projection coefficient \mathcal{N} in equation (37). This subsection advances that problem by showing that the *functional form* of the projection (specifically, that the leading curvature term is $\int R\sqrt{-g} d^4x$ and not $\int R_{\mu\nu}R^{\mu\nu}\sqrt{-g} d^4x$ or any other combination) is forced by a universal result from differential geometry: the Seeley–DeWitt heat-kernel expansion [55, 57]. What remains open (the numerical coefficient, which fixes G in terms of r_0 , c , and d_0) is a separate calculation from what this subsection resolves.

Setup. The projection operator Π is characterised by a kernel $K: \mathcal{C} \times M \rightarrow \mathbb{R}$ satisfying

$$\Pi(A)[x] = \int_{\mathcal{C}} A[\psi] K(\mathcal{C}, x) d\mu(\mathcal{C}) \quad (40)$$

for any functional A on configuration space. The kernel inherits the permutation invariance of the Bernoulli measure μ ; its spatial covariance is controlled by the substrate coherence length d_0 , so that $K(\mathcal{C}, x)$ couples configurations within a neighbourhood of radius $\sim d_0$ of the pre-image $\rho^{-1}(x)$.

Vacuum decomposition. At the vacuum $\psi = r_0 e^{i\theta}$, the configuration-space gradient decomposes as

$$|\nabla_{\mathcal{C}}\psi|^2 = r_0^2 |\nabla_{\mathcal{C}}\theta|^2 + |\nabla_{\mathcal{C}}\sigma|^2 + \mathcal{O}(\sigma^2) \quad (41)$$

where $\sigma = |\psi| - r_0$ is the amplitude (metric-fluctuation) mode and θ is the Goldstone (phase) mode. The Goldstone term $r_0^2 |\nabla_{\mathcal{C}}\theta|^2$ projects to the gauge and matter sector (Sections 7 and 8). The amplitude term $|\nabla_{\mathcal{C}}\sigma|^2$ encodes how the instantiated geometry responds to deformations of the realisation map ρ : a deformation $\delta\rho$ induces a metric variation $\delta g_{\mu\nu} = \mathcal{L}_{\delta\rho} g_{\mu\nu}$ (the Lie derivative), and $\nabla_{\mathcal{C}}\sigma$ in the direction of $\delta\rho$ is the corresponding functional derivative of the amplitude. It is this term that, when projected via Π , produces the gravitational action.

The heat-kernel form of the projected integral. The projected amplitude gradient can be written as a smeared trace over M :

$$\Pi(c|\nabla_{\mathcal{C}}\sigma|^2) |_{\sigma=0} = c r_0^2 \int_M \tilde{K}(x, x; d_0^2) \sqrt{-g} d^4x \quad (42)$$

where $\tilde{K}(x, y; t)$ is the coincidence limit of the projection kernel restricted to metric fluctuations, with $t = d_0^2$ playing the role of the “heat-kernel time.” The Seeley–DeWitt theorem requires \tilde{K} to be, to leading order, the heat kernel of a *Laplace-type* operator $P = -(g^{\mu\nu} \partial_{\mu} \partial_{\nu} + \dots)$ whose principal symbol is the spacetime metric g . This is not a consequence of smoothness, positivity, symmetry, and localisation: a normalized geodesic Gaussian satisfies all four properties yet has diagonal curvature coefficient $\frac{1}{3}R$ rather than $\frac{1}{6}R$, as a direct calculation on S^2 confirms. The $\frac{1}{6}$ is a property of the heat kernel of the Laplacian specifically, not of the broad class of localized kernels.

The Laplace-type condition is formalised as Postulate 4 (Section 1.7): \tilde{K} is defined as the heat kernel of the projected Boolean Laplacian $\Pi(L)$, with the metric g defined as the principal symbol of $\Pi(L)$. This breaks the potential circularity: g is not presupposed when defining \tilde{K} ; it is read off from the operator. With this identification the Seeley–DeWitt expansion applies and the universal coefficient $a_1 = \frac{1}{6}$ forces the leading curvature term to be $\int R \sqrt{-g} d^4x$.

Taken as an independent postulate, P4 would be non-trivial: stipulating that $\Pi(L)$ is Laplace-type is close to postulating the kinematic skeleton of a metric field theory directly. The Mosco-convergence argument of the following paragraph derives this condition from P1–P3, reducing P4 to a theorem. The honest reading of PST’s gravitational sector is therefore that

General Relativity is recovered from three foundational postulates, with the operator-type condition following as a structural consequence of Bernoulli isotropy.

Deriving the Laplace-type condition from the Bernoulli measure (Mosco convergence). The conditional character of P4 can be substantially reduced by an argument using Mosco convergence of Dirichlet forms [92]. Define the *Boolean Dirichlet form* on $L^2(\mathcal{P}(D), \mu)$ by

$$E^B(u, v) = \sum_{a \in D} \int_{\mathcal{P}(D)} (\partial_a^B u)(C) (\partial_a^B v)(C) d\mu(C), \quad (43)$$

where ∂_a^B is the Boolean difference operator of equation (17) and the generator of E^B is the Boolean Laplacian L^B . For smooth functions $f, h \in C^\infty(M)$ pulled back via Π^\dagger , the isotropy of the Bernoulli measure μ forces the covariance matrix of the displacement $\rho(a)^\mu - x^\mu$ to be proportional to $g^{\mu\nu}(x)$:

$$\frac{1}{|D|} \sum_{a \in D} (\rho(a)^\mu - x^\mu) (\rho(a)^\nu - x^\nu) \longrightarrow \frac{d_0^2}{4} g^{\mu\nu}(x) \quad (44)$$

(no preferred direction exists in the substrate, so the only $O(4)$ -covariant rank-2 tensor built from g is $g^{\mu\nu}$ itself). It follows that

$$\frac{1}{|D|} E^B(\Pi^\dagger(\Pi f), \Pi^\dagger(\Pi h)) \xrightarrow[|D| \rightarrow \infty]{\text{Mosco}} \frac{d_0^2}{4} \int_M g^{\mu\nu} \partial_\mu f \partial_\nu h \sqrt{-g} d^4x =: \frac{d_0^2}{4} E^{\text{Riem}}(f, h), \quad (45)$$

where E^{Riem} is the Riemannian Dirichlet form whose generator is the Laplace–Beltrami operator Δ_g . By Mosco’s theorem [92], this Γ -convergence of quadratic forms implies strong resolvent convergence of their generators:

$$(|D|^{-1} L^B + \lambda)^{-1} \Pi^\dagger \xrightarrow{|D| \rightarrow \infty} \Pi(\Delta_g + \lambda)^{-1} \quad \forall \lambda > 0. \quad (46)$$

The generator of the limiting form is Δ_g , a Laplace-type operator with principal symbol $g^{\mu\nu}$. Therefore $\Pi(L^B)$ is Laplace-type in the scaling limit, and P4 (the assumption that the projection kernel is the heat kernel of a Laplace-type operator) follows from P1 (the distinction primitive, with the Bernoulli measure derived in Section 3), P2 (asymmetric tension), and P3 (the Landau-Ginzburg modal potential). This residual is thereby resolved within the scaling argument.

What is assumed at the Mosco step. The argument above derives the *Laplace-type* character of $\Pi(L^B)$ in the scaling limit; it does *not* derive the dimensionality $n = 4$ or the Lorentzian signature $(-, +, +, +)$ from S_D -symmetry of the Boolean substrate alone. Three structural elements enter at the realisation step beyond P1–P3: (i) the realisation map $\rho : D \rightarrow M$ presumes the target manifold M is smooth and finite-dimensional; (ii) the dimensionality $n = 4$ enters via the index structure $\rho(a)^\mu - x^\mu$ ($\mu = 0, 1, 2, 3$) and the notation $O(4)$ in the isotropy step, S_D -symmetry on the discrete substrate by itself does not specify the target dimension, and the natural large- $|D|$ limit of the Hamming cube $\{0, 1\}^{|D|}$ is an infinite-dimensional Gaussian

object unless an additional structural element reduces it; (iii) the Lorentzian signature is the sign cut at the null threshold τ_{null} in equation (29). Under the native reading these three are not free posits but follow from the verified Connes-spectral-triple structure of the substrate. Computation 3 shows that on the parity-selected sub-limit (odd $|D|$) the substrate's complementation-and-parity real structure makes it a real Connes spectral triple of KO-dimension six with the Lorentzian-compatible sign triple $(+, +, -)$. Computation 4 shows that the Mosco-limit spacetime $M = \mathbb{R} \times S^3$ has a unique spin structure (KO-dimension four). KO additivity then forces the total to $4 + 6 = 10 \equiv 2 \pmod{8}$, the Connes Lorentzian Standard-Model value. This is not an assumption but the verified arithmetic of the product spectral triple; $n = 4$ is forced because any other value would either contradict the substrate's parity-verified KO-6 or fail KO additivity to ten. Lorentzian signature follows by Computation 2 hyperbolic well-posedness of the Goldstone wave equation (Section 14.2); the null threshold τ_{null} is the sign cut that hyperbolicity requires, not an independent input. PST therefore derives the dimension and signature of M from the three postulates plus the verified spectral-triple structure of the substrate and the Mosco-limit spacetime. Deriving these directly from S_D -symmetry alone, without invoking the Connes spectral-triple machinery, would be a deeper but currently unattained result.

Lorentzian signature from the directionality of modal sublimation. The argument above yields Euclidean signature in the $|D| \rightarrow \infty$ limit. Lorentzian signature $(-+++)$ arises from the directionality of modal sublimation: the tension $T(C)$ exceeds the threshold τ in a preferred configuration direction that, after instantiation, becomes the time-like direction. This breaks the $O(4)$ symmetry of the Euclidean scaling limit to the Lorentz group $O(1, 3)$. The Wick rotation is therefore a structural consequence of P2 (asymmetric tension) together with the threshold-crossing of P3, not an additional postulate. The Computation 2 hyperbolic well-posedness argument (Section 14.2) gives the same conclusion by a different route, showing it is the unique signature compatible with a well-posed Cauchy problem for the Goldstone wave equation.

The Seeley–DeWitt expansion. For any smooth kernel of this type on a four-dimensional (pseudo-)Riemannian manifold, the asymptotic expansion as $d_0 \rightarrow 0$ gives [55, 56, 57]:

$$\tilde{K}(x, x; d_0^2) = \frac{\mathcal{N}}{(4\pi d_0^2)^2} \left[1 + \frac{d_0^2}{6} R(x) + d_0^4 \mathcal{A}_2(x) + \mathcal{O}(d_0^6) \right] \quad (47)$$

where \mathcal{N} is the normalisation of the kernel (fixed by $\int_M \tilde{K} \sqrt{-g} d^4x = 1$ on flat space), the coefficient $\frac{1}{6}$ in front of R is the *universal* first Seeley–DeWitt coefficient a_1 for a scalar field, and $\mathcal{A}_2(x) = \frac{1}{180}(R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} - R_{\mu\nu}R^{\mu\nu} + \frac{1}{5}\nabla^2 R) + \frac{1}{72}R^2$ is the second coefficient. The universality of $\frac{1}{6}$ is the key: it depends only on the spacetime dimension and the spin of the projected field, not on any property of the PST kernel K .

Integration and the Einstein-Hilbert structure. Substituting equation (47) into equation (42):

$$\Pi(c|\nabla_c\psi|^2)|_{\psi=r_0} = \frac{c r_0^2 \mathcal{N}}{(4\pi d_0^2)^2} \int_M \sqrt{-g} d^4x + \frac{c r_0^2 \mathcal{N}}{6(4\pi)^2 d_0^2} \int_M R \sqrt{-g} d^4x + \mathcal{O}(d_0^4 R_{\mu\nu\rho\sigma}^2) \quad (48)$$

The first term is a constant (vacuum energy) that is absorbed into the cosmological constant. The second term has the precise structure of the Einstein-Hilbert action, confirming equation (37) to leading curvature order. The higher-order terms are suppressed by $(k d_0)^2$ at momenta $k \ll 1/d_0$, consistent with the assertion in equation (37).

The coefficient and Newton's constant. Comparing the coefficient of $\int R\sqrt{-g}d^4x$ in equation (48) with $1/(16\pi G)$:

$$\frac{c r_0^2 \mathcal{N}}{6(4\pi)^2 d_0^2} = \frac{1}{16\pi G} \quad \implies \quad G = \frac{6(4\pi)^2 d_0^2}{16\pi c r_0^2 \mathcal{N}} = \frac{6\pi d_0^2}{c r_0^2 \mathcal{N}} \quad (49)$$

The quantity \mathcal{N} is the normalisation integral of the PST kernel K over the Bernoulli measure:

$$\mathcal{N} = \int_{\mathcal{C}} K(\mathcal{C}, x) d\mu(\mathcal{C}) \quad (50)$$

evaluated at any point $x \in M$ (the result is independent of x by permutation invariance of μ). This is the kernel integral that determines the numerical coefficient \mathcal{N} . Equation (49) shows that once \mathcal{N} is computed from the Bernoulli measure, the gravitational constant G is fully determined in terms of c , r_0 , and d_0 .

Summary. The Seeley–DeWitt argument establishes three things without additional assumption. First: the functional form of the gradient projection is $\int R\sqrt{-g}d^4x$ at leading curvature order, forced by the universal coefficient $a_1 = 1/6$. Second: the higher-order corrections are $R_{\mu\nu\rho\sigma}^2$ terms suppressed by $(k d_0)^2$. Third: G is determined by \mathcal{N} via equation (49). The value of \mathcal{N} is computed in the following subsection.

7.7. Dirac-operator convergence: structural prerequisites and the residual analytic estimate

The Mosco argument above establishes strong resolvent convergence of the Boolean Laplacian to Δ_g ; the Connes spectral-triple framework assumption requires the stronger statement that the rescaled *substrate* spectral triple converges to the canonical Dirac triple

$$(C^\infty(\mathbb{R} \times S^3), L^2(\mathcal{S}), D_{\mathbb{R} \times S^3}, J, \gamma)$$

in an appropriate substrate-to-continuum metric. The original target was Latrémolière's spectral propinquity [98, 99] (operator-level convergence of Dirac data); the actual closure target for PST is the relaxed norm-equivalence criterion of equation (55) below (§14.2 the Dirac-aware lift), in the distance framework of [97]. The relaxation is structurally forced by the foundational object's fixed internal Hilbert space $H_F = \mathbb{C}^8$ (§8.8); details below.

Computation 9 (`computation_09.py`) decomposes this convergence into three ingredients and reports the status of each. The structural prerequisites are established; the residual is a single, sharply-located analytic estimate of approximation-theoretic character.

Three ingredients.

1. *Mosco / strong-resolvent convergence of the Boolean Dirichlet forms to the S^3 Laplace–Beltrami form.* Established in §7.6; this is the form-level convergence on which the substrate’s spectral-triple identification rests.
2. *Unique spin structure on $M = \mathbb{R} \times S^3$.* The second Stiefel–Whitney class $w_2(M)$ vanishes and $H^1(S^3; \mathbb{Z}_2) = 0$, so M admits a unique spin structure and the canonical Dirac operator D_M is uniquely determined (Computation 4, `computation_04.py`). Friedrich’s eigenvalue formula [94, 95] gives the S^3 Dirac spectrum $\pm(n + \frac{3}{2})$ with multiplicity $2(n + 1)(n + 2)$.
3. *Boolean–CAR Clifford-action convergence.* This is the residual: the discrete CAR Clifford action must converge to the continuum Clifford action so that the *first-order* operators are compatible at the substrate-to-continuum metric, not merely their squares. The original natural framework is Latrémolière’s spectral propinquity on the moduli of spectral triples; the actual PST closure target is the relaxed norm-equivalence form below (equation (55); §14.2 the Dirac-aware lift).

Algebraic and spectral prerequisites of ingredient 3 are verified. At every finite substrate size D the Boolean–CAR generators χ_a, χ'_a ($a = 1, \dots, D$) satisfy the $\text{Cl}(0, 2D)$ anticommutation $\{\chi_a, \chi_b\} = 2\delta_{ab} = \{\chi'_a, \chi'_b\}$, $\{\chi_a, \chi'_b\} = 0$ *exactly* (Computation 9 §2, machine-precision check up to $D = 6$); the Walsh-mode binomial multiplicities $\binom{D}{k}$ flow under the LG-induced reweighting (Sections 3 and 6) to the $(k + 1)^2$ spherical-harmonic multiplicities at every k in the thermodynamic limit (Computation 9 §3, exhibited at $D = 6, 8, 10, 12$); the bridge map b_n from Walsh-truncation polynomials to spherical-harmonic polynomials is constructed explicitly (Computation 9 §7); and the low-mode commutator gap vanishes *exactly* on Walsh-weight ≤ 1 test functions at every D ($\|[\chi_a^{(D)}, f] - b([\gamma_a, f])\| = 0$, verified numerically for $D = 3, 4, 5, 6$ in Computation 9 §8).

The residual estimate, structurally. The full first-order convergence reduces, given the four items above, to a uniform-in- D Sobolev-tail estimate on the dense polynomial subalgebra of $C^\infty(S^3)$:

$$\|[\chi_a^{(D)}, f_D] - b_D([\gamma_a, f])\|_{\text{op}} \leq C D^{-s/2} \|f\|_{H^{s+1}(S^3)}, \quad f \in H^{s+1}(S^3), \quad (51)$$

where $\chi_a^{(D)}$ is the discrete CAR Clifford generator at substrate size D , f_D is the Walsh-truncation of f at weight $k_D = D^{1/2}$, b_D is the bridge map of Computation 9 §7, and γ_a is the continuum Clifford action. Computation 9 §11 gives a six-step proof structure for this estimate. Steps 1–4 (spherical-harmonic decomposition, Walsh truncation, tail-norm bound,

Clifford-derivation lift) and step 6 (triangle inequality + Walsh–Hadamard unitarity) are standard Sobolev / Clifford-derivation manipulations on compact spin manifolds; the numerical check in Computation 9 §13 confirms the tail-rate scaling. Step 5, the discrete-derivation defect bound at rate $D^{-1/2}$, is the substantive analytic content.

What step 5 requires. A direct inspection of Latrémolière 2018 [98, 99] (Computation 9 §15) confirms that step 5 is *not* closed by routine citation under the original spectral-proximity framing: spectral-proximity convergence is established case-by-case for each finite-dimensional approximation scheme. The two sphere-approximation precedents in the literature (Rieffel’s matrix-algebra $\rightarrow S^2$ convergence (Berezin/Kähler scheme, in the quantum Gromov–Hausdorff distance, [96]) and Bhattacharyya–Singla’s Toeplitz-algebra $\rightarrow S^{2k+1}$ convergence (Bergman scheme, also in the quantum Gromov–Hausdorff distance, [100])) both use complex-projective / coherent-state machinery that does not directly port to the Boolean cube’s real Dirichlet-form structure. Diaconis–Shahshahani’s hypercube-random-walk convergence theorem [101] gives the appropriate measure-level convergence rate $(\log D)/D$; lifting this to operator-level convergence in the original spectral-proximity framework is not the route PST commits to. Instead, the actual closure target (§14.2 the Dirac-aware lift) is the relaxed norm-equivalence criterion of equation (55), for which the Fréchet smooth Lip-norm machinery from L -round closure (Computations 20–23, *section round-trip* paragraph above) plausibly transfers; the empirical $\gamma_D = 0.007$ at $D = 8$, $N = 20$ at the Walsh-weight cutoff (Computation 36) is the finite-truncation verification, with the analytical closure now achieved via the refined symmetric-monomial bridge (Computations 39–41).

Status. The Connes real spectral-triple identification has six checkable pieces; five of them (Mosco resolvent convergence of the Boolean Laplacian (§7.6), unique spin structure (Computation 4), exact discrete Clifford algebra at every D (Computation 9 §2), exact multiplicity flow (Computation 9 §3), exact low-mode commutator gap (Computation 9 §8)) are verified at the structural level with no obstruction. The sixth piece, the Sobolev-tail bound of equation (51), is the residual. The residual splits cleanly into two sub-statements (an observation absent from the earlier drafts, see [97]): a quantum Gromov–Hausdorff section round-trip lemma analogous to Bhattacharyya–Singla 2022 Lemma 3.2; and the lift of that round-trip from quantum Gromov–Hausdorff to spectral proximity, the genuinely Dirac-aware statement of equation (51).

Section round-trip in quantum Gromov–Hausdorff distance, numerically closed.

Let $\pi_D : A_D \rightarrow C(S^3)$ be Walsh-coefficient truncation at weight $|S| \leq k_D := \lfloor \sqrt{D} \rfloor$, and let $\sigma_D : C(S^3) \rightarrow A_D$ be the trivial section (the same low-weight content viewed in A_D). The PST analog of [100] Lemma 3.2 asks for a sequence $\gamma_D \rightarrow 0$ and a Lip-norm L_D on A_D with

$$\|T - \sigma_D(\pi_D(T))\|_{\text{op}} \leq \gamma_D \cdot L_D(T) \quad \text{for all } T \in A_D. \quad (52)$$

A numerical sweep across $D = 4, 6, \dots, 20$ and three test classes (single Walsh modes; constant-coefficient Walsh tail; random Walsh coefficients) tested eleven candidate Lip-norms (polynomial sup, polynomial Sobolev, exponential sup, exponential Sobolev) under equation (52) (Computation 20). The result is decisive in two halves.

Polynomial Walsh-weight Lip-norms FAIL. For every polynomial weight $w(|S|) = (1 + |S|)^\alpha$ ($\alpha = 1, 2, 3$ or the corresponding Sobolev H^p , $p = 1, 2$), the constant-coefficient tail $T = \sum_{|S| > k_D} \chi_S$ has $\|T_{\text{tail}}\|_{\text{op}} = 2^D$ (the count of tail subsets, all evaluating to +1 at the all-zero state), while $L_D(T) \sim D^\alpha$ grows only polynomially. Hence $\gamma_D \geq 2^D/D^\alpha$ blows up exponentially with D , and (52) fails for every polynomial Lip-norm. Empirical log-log slopes are +6.6 for linear weights down to +2.5 for H^2 Sobolev, all positive.

The Fréchet smooth Lip-norm SUCCEEDS. Define the Fréchet smooth Lip-norm

$$L_F(T) := \sup_{j \geq 0} \frac{\|\text{ad}_D^j(T)\|_{\text{op}}}{j!}, \quad \text{ad}_D(T) := [D, T], \quad D = \sum_a \chi_a, \quad (53)$$

the standard smooth-subalgebra Lip-norm of the spectral triple. Higher-order Dirac commutators of single Walsh modes satisfy $\|\text{ad}_D^k(\chi_S)\|^{1/k} \rightarrow 2\sqrt{D}$ at large k (independent of $|S|$; Computation 22), so by Stirling $L_F(\chi_S) \sim e^{2\sqrt{D}}$. *The single-mode bound is now rigorous* (Computation 50): the Jordan–Wigner Clifford structure gives $D^2 = D \cdot I$ on $(\mathbb{C}^2)^{\otimes D}$, hence $\|D\|_{\text{op}} = \sqrt{D}$, and the standard left-right multiplication estimate yields $\|\text{ad}_D^j(\chi_S)\|_{\text{op}} \leq (2\sqrt{D})^j \|\chi_S\|_{\text{op}}$ uniformly in j, S, D . Combined with Stirling this proves the single-mode L-round closure $\|\chi_S\|_{\text{op}}/L_F(\chi_S) \leq e^{-2\sqrt{D}} \rightarrow 0$ as a theorem, not merely a numerical observation. Direct measurement on the constant-coefficient tail (Computation 23) yields $\gamma_D \approx 0.37 \cdot e^{-0.28D}$ across $D = 4, 6, 8, 10, 12$, confirming exponential decay. The tail-sum closure now reduces to an explicit linear-algebra estimate via the block-decomposition identity: because $D^2 = D \cdot I$ on $(\mathbb{C}^2)^{\otimes D}$, the spectrum of D is $\{\pm\sqrt{D}\}$ with spectral projectors $P_\pm = (I \pm D/\sqrt{D})/2$, and direct calculation shows ad_D acts as 0 on the diagonal blocks T_{++}, T_{--} and as $\pm 2\sqrt{D}$ on the off-diagonal blocks $T_{\pm\mp} := P_\pm T P_\mp$. The block-anti-diagonal structure of $T_{+-} - T_{-+}$ then gives, for every T and every $j \geq 1$,

$$\|\text{ad}_D^j(T)\|_{\text{op}} = (2\sqrt{D})^j \cdot \max(\|T_{+-}\|_{\text{op}}, \|T_{-+}\|_{\text{op}}), \quad (54)$$

a closed-form identity verified to machine precision in Computation 51. The Fréchet smooth Lip-norm therefore admits the exact factorisation $L_F(T) = \max(\|T_{+-}\|, \|T_{-+}\|) \cdot M(D)$ with $M(D) := \sup_{j \geq 1} (2\sqrt{D})^j / j!$, and the L-round rate becomes $\gamma_D = \|T_{\text{tail}}\|_{\text{op}} / (\max(\|T_{+-}\|, \|T_{-+}\|) \cdot M(D))$. Lemma 2 (tail-sum L_F scaling, in the formulation of `research/dirac_proof.md`) is thereby *reduced* from the open interference question of higher-order commutators to the explicit operator-norm bound on the structured off-diagonal block $P_+ T_{\text{tail}} P_-$. Numerically, $\max(\|(T_{\text{tail}})_{+-}\|, \|(T_{\text{tail}})_{-+}\|) = 2^{D-1} - D$ exactly across $D \in \{4, 5, 6, 7, 8\}$ at cutoff $k_D = 2$, yielding the true asymptotic rate $\gamma_D = O(D^{1/4} e^{-2\sqrt{D}})$ (the pre-asymptotic empirical exponential fit $0.37 \cdot e^{-0.28D}$ of Computation 23 crosses over to the square-root rate near $D \approx 51$).

Either rate satisfies $\gamma_D \rightarrow 0$, so `L`-round closes. The remaining analytical step is the rigorous proof of the off-diagonal block formula. Combined with the bridge construction of [97] (Definition 5.1, Theorem 5.2), this gives convergence of $(\text{Lip}(A_D), L_F) \rightarrow (C(S^3), L)$ in the quantum Gromov–Hausdorff distance at exponential rate in substrate size.

Why polynomial Lip-norms fail and the Fréchet smooth one succeeds is structural. The Walsh substrate has 2^D modes per substrate size D , whereas the Bergman basis in [100] has only polynomially many. The Lip-norm must dominate the mode count, hence must grow at least exponentially in the Walsh weight. The Fréchet smooth Lip-norm does this automatically: every higher-order Dirac commutator contributes, and the supremum over j of $\|\text{ad}_D^j(T)\|/j!$ captures the infinite-order Sobolev structure of the smooth subalgebra. So PST’s `L`-round closes within the standard Connes–Chamseddine spectral-triple framework, with no artificial weight choice.

Dirac-aware lift to substrate-to-continuum metric. With the section round-trip numerically closed, the residual is the lift of the quantum-Gromov–Hausdorff convergence (Lip-norm only) to a Dirac-aware substrate-to-continuum metric. The original target was Latrémolière’s spectral propinquity ([98, 99]) at the level of equation (51), requiring full operator equality of bridge commutators. The actual closure target for PST (below) is the relaxed norm-equivalence criterion equation (55) in the [97] distance framework, strictly weaker than spectral propinquity but PST-appropriate as explained below. Either target requires the spinor extension of the bridge map b_D implemented in Computation 32.

The discrete-side audit (Computation 21, Computation 24) confirms the algebraic structure: the single-commutator Dirac Lip-norm satisfies $\|[D, \chi_S]\|_{\text{op}} = 2\sqrt{|S|}$ exactly, matching the Friedrich Dirac eigenvalue scaling $\sqrt{l(l+2)} \approx \sqrt{l}$ on round S^3 at degree $l = |S|$ ([94]; tabulated and cross-checked against Friedrich’s eigenspinor decomposition in Computation 25); and $[\chi_a, \chi_S]$ is Hilbert–Schmidt orthogonal to the entire Walsh subalgebra at every D and every weight. The bridge defect therefore lives entirely in the non-abelian Pauli sector of the discrete CAR algebra, perpendicular to the abelian Walsh sector.

The bridge construction was narrowed by two structural results. Computation 26 establishes a structural no-go for the entire bijective single-mode class: the discrete-side commutator op-norm is exactly $\|[\chi_a^{\text{Cliff}}, \chi_S]\|_{\text{op}} = 2$ (independent of D and of $|S|$), while the continuum-side norm scales as $\sqrt{l(l+1)}$ via the $\text{SU}(2)$ Casimir; the `L`-round scaling on the same bridge fixes the overall weight-graded scalar, leaving no internal freedom to reconcile the two scales. Computation 27 then implements the full BS 2022 Toeplitz framework on $H_\alpha^2(B^2)$ explicitly. The Toeplitz commutator $[T_{z_1^k}, T_{z_1^k}]$ is diagonal on the monomial basis with an explicit entry $\Lambda(J, k, \alpha)$; at $k = 1$ its operator norm equals $1/(\alpha + 3)$ to machine precision (verified against the analytical formula). The Toeplitz operator norm $\|T_{z_1^k}\|_{\text{op}, \alpha}$ approaches 1 and is α -independent at leading order. The ratio $R(k, \alpha) = \|\text{comm}\|/\|T_{z_1^k}\|$ is therefore α -DEPENDENT, providing precisely the non-proportional α -scaling needed. The BS 2022 closure mechanism is thus EXPLICITLY VERIFIED on $H_\alpha^2(B^2)$: $\alpha = \alpha(D)$ supplies the internal degree of freedom

needed to satisfy both fidelity and gap rate.

With the BS 2022 Toeplitz framework on $H_\alpha^2(B^2)$ now explicitly verified (Computation 27), the explicit bridge identifying the PST substrate $\text{Cl}(0, 2D)$ algebra with the Toeplitz algebra was constructed in successive refinements: the multiplicative bridge of Computation 28 documented spurious cross-commutators between mixed holomorphic and anti-holomorphic factors; restriction to the holomorphic Toeplitz subalgebra (Computation 29) restored the abelian Walsh structure at the cost of fidelity decay with weight; the optimal- α diagnostic (Computation 30) showed that no single scalar α can match the bimodal substrate commutator $\{0, 2\}$ since the Toeplitz commutator is a continuous function of α ; and the tensor spinor extensions (Computation 31) decouple in the operator norm and therefore do not help. Computation 32 implements the non-tensor coupling identified as missing: the round- S^3 Dirac $D = \sum_a \sigma_a \otimes J_a$ on $H_\alpha^2(B^2) \otimes \mathbb{C}^2$, where J_a are the $\text{SU}(2)$ generators acting on the Bergman monomial basis as raising/lowering operators. The framework is operational (Hermitian Dirac verified, $\text{su}(2)$ algebra to machine precision, eigenvalue range approximating Friedrich's $\pm(n + 3/2)$ S^3 spectrum) and the L-comm commutator $\|[D, b_{\text{holo}}(\chi_S) \otimes I]\|_{\text{op}}$ is non-zero with normalised ratios 1.02, 1.83–2.29, 2.45–3.04 at $|S| = 1, 2, 3$, approaching the substrate target $2\sqrt{|S|}$. Mosco averaging across same-weight Walsh modes combined with α -rescaling (Computation 33) produces ratios that monotonically approach 1 with weight (0.60, 0.71, 0.87 at $k = 1, 2, 3$), and scaling to $D = 6$, $N = 12$ pushes the $k = 3$ ratio to 0.96 within 4% of target (Computation 34), with a clean structural crossover at $k \sim 3$ reflecting the $3 + 1$ structure of the round S^3 Dirac.

These successes measure $\|\text{bridge commutator}\|$ against the substrate norm $2\sqrt{|S|}$, however, which is weaker than the genuine L-comm criterion. The genuine criterion is the *homomorphism gap* $\|\text{bridge}([D_{\text{sub}}, \chi_S]) - [D_\alpha, \text{bridge}(\chi_S)]\|_{\text{op}}$, tested in Computation 35 with the natural multiplicative extension $\text{bridge}(\chi_a^{\text{Cliff}}) := J_a \otimes \sigma_a$. The result is structurally negative: the homomorphism gap is large ($\text{gap}/\|\text{LHS}\| \in [0.68, 1.01]$) and does NOT shrink with (D, N) , uniformly across $(4, 5)$, $(4, 10)$, $(6, 12)$. The diagnosis is a dimensional mismatch. The substrate has D anti-commuting Clifford generators χ_a^{Cliff} while the S^3 spinor fibre \mathbb{C}^2 carries only three Pauli matrices; the cyclic assignment forced upon $D > 3$ identifies $\text{bridge}(\chi_a^{\text{Cliff}}) = \text{bridge}(\chi_b^{\text{Cliff}})$ for $a \equiv b \pmod{3}$, which is incompatible with a homomorphism since the substrate generators anti-commute.

The L-comm closure target is therefore the *relaxed norm-equivalence criterion*:

$$\left| \|[D_\alpha, \text{bridge}(\chi_S)]\|_{\text{op}} - \|[D_{\text{sub}}, \chi_S]\|_{\text{op}} \right| \leq \gamma_D \cdot \|\chi_S\|_{\text{op}}, \quad \gamma_D \rightarrow 0 \text{ at exponential rate, } (55)$$

on the Walsh-weight-bounded subalgebra $|S| \leq k_D = \lfloor \sqrt{D} \rfloor$. This is strictly weaker than the original spectral-proximity criterion ([98, 99]) but is the appropriate target for PST. The choice is structurally forced: the alternative – enlarging the bridge spinor fibre to \mathbb{C}^{2^D} to make the operator equality $\text{bridge}([D_{\text{sub}}, \chi_S]) = [D_\alpha, \text{bridge}(\chi_S)]$ achievable – is incompatible with

the foundational object's fixed internal Hilbert space $H_F = \mathbb{C}^8$ (§8.8). The spectral-triple framework insists on a single Hilbert space per triple, and an `L_comm`-specific \mathbb{C}^{2^D} spinor structure that grows with the substrate size has no home alongside the foundational \mathbb{C}^8 .

Why the relaxed criterion suffices for PST. Every load-bearing derivation in the paper uses the substrate's spectral-triple structure at finite D : Newton's G via the spectral action $\text{Tr } f(D/\Lambda)$ (§7.9), the gauge group $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$ via the unimodular unitaries of A_F (§8), gauge bosons and Higgs as inner fluctuations of D (§8), the KO-dimension assignment (Computation 3). In every case the continuum limit is the *endpoint* of the construction, not the input. The spectral action depends on the spectral data (eigenvalues and traces) of D , which the relaxed criterion preserves under equation (55); full operator equality of $\text{bridge}([D_{\text{sub}}, \chi_S])$ with $[D_\alpha, \text{bridge}(\chi_S)]$ is not invoked. Closure under (55) therefore suffices for the Standard Model derivation chain.

The distance framework is [97]'s, the same machinery already used to lift `L_round` closure to quantum Gromov–Hausdorff convergence. Computations 33, 40, 42 establish equation (55) at *finite- N matched truncation* (Bergman dimension $N \sim 2.5 D$): at the Walsh-weight cutoff $k_D = \lfloor \sqrt{D} \rfloor$, the gap $\gamma_D = |1 - \text{ratio}|$ at $N = 20$ reaches $\gamma_D = 0.007$ at $D = 8$ (within 0.7% of the substrate target), with empirical fit $\gamma_D \sim 19.3 \cdot e^{-0.93D}$ across $D = 4, 6, 8$.

Caveat on the rate. Computations 37–38 isolate the *infinite*-Bergman-truncation contribution to γ_D via a closed-form sup-norm criterion on ∂B^2 (Section 14.2, the Dirac-aware lift). The infinite- N analytical gap at $k_D = 2$ across $D = 4, 6, 8$ is $\gamma_D^{(\infty)} = 0.389, 0.193, 0.048$, fitting an exponential at rate $c \approx 0.527$, not 0.93. The remainder of the apparent finite- $N = 20$ rate ($0.93 - 0.527 \approx 0.4$) is contributed by the *truncation correction* $\gamma_D^{(\infty)} - \gamma_D^{(N)}$, which is specific to the matched truncation $N(D) \sim 2.5 D$. In addition, at the cutoff transition $k_D = 2 \rightarrow 3$ (between $D = 9$ and $D = 10$), the analytical gap jumps from 0.048 to 0.444 (Computation 38 at $D = 10$), so the infinite- N closure is not uniform in D across cutoff transitions. The exponential closure verified in Computation 36 is therefore best read as a *matched-truncation L_comm criterion at the working $N(D) \sim 2.5 D$ scaling*, not as a spectral-proximity statement on the infinite Bergman space.

Refined bridge – demonstrated closure. The infinite- N closure of equation (55) requires a bridge construction whose substrate-to-Bergman ratio matches $2\sqrt{k}$ uniformly across $k \leq k_D$, smoothing the cutoff discontinuity exposed by Computation 38. Computations 39–41 construct the *symmetric-monomial* bridge $b_C(\chi_S) := T_{f,|S|}$ with

$$f_k(z_0, z_1) = (z_0 z_1)^{m(k)-1} + \beta(k) (z_0 z_1)^{m(k)} + \alpha(k) (z_0 z_1)^{m(k)+1}.$$

The bare $m = 1$ monomial gives ratio = 2 = $2\sqrt{1}$ exactly at $k = 1$; two-monomial superposition (Computation 40) gives gap below 10^{-3} at all even k ; three-monomial superposition (Computations 41–42) closes the residual gap at odd k and extends the closure uniformly across $k = 1, 2, \dots, 12$, with maximum gap 3.6×10^{-4} at $k = 9$ (several k close exactly to grid precision,

including $k = 1, 3, 5, 12$). The refined bridge gives $\gamma(k)$ that is *D-independent*, eliminating the cutoff discontinuity. The Mosco averaging becomes trivial (all Walsh modes of weight k map to the same operator), so the substrate–Bergman correspondence is via a k -indexed symmetric construction rather than the multiplicative site-by-site map of Computations 32–36.

Remaining analytical work. Promote the numerical closure of Computations 39–41 to a rigorous proof: (i) verify the Berezin-transform compatibility of the refined bridge with the weighted-Bergman spectral-triple structure ([100]); (ii) establish the spectral-action invariance under the refined bridge; (iii) confirm that the D -independent $\gamma_\infty(k) \rightarrow 0$ result extends to a uniform bound across $k \leq k_D$ for all D . The analytical framework (closed-form Leibniz identity $[J_a, T_F] = T_{J_a F}$ on holomorphic Toeplitz symbols, corrected pointwise op-norm of the matrix-valued symbol over ∂B^2 , sup-norm criterion validated numerically to four digits) is now set up; the remaining work is the rigorous proof. Estimated effort: research-paper-level operator-algebras analysis (1–2 months).

The section round-trip and the verifications above are unaffected. This is the restated content of §14.2, the Dirac-convergence item; see also Appendix C rows 22–41.

7.8. Kernel normalisation: the finite-internal-factor trace

Equation (49) expresses Newton’s constant as $G = 6\pi d_0^2 / (c r_0^2 \mathcal{N})$, with \mathcal{N} the kernel normalisation. In the spectral-triple reading \mathcal{N} is the trace over the finite internal factor F ,

$$\mathcal{N} = \text{Tr}_{\mathcal{H}_F}(\mathbf{1}), \tag{56}$$

the fermionic multiplicity of the internal Hilbert space of one generation. This is a finite number fixed by the matter content, not a free parameter; it enters the spectral-action coefficient $1/G \sim f_2 \Lambda^2 \mathcal{N}$ alongside the cutoff-function moment f_2 .

Closed form of Newton’s constant. Absorbing the fixed multiplicity \mathcal{N} and the moment f_2 into the gradient coupling c , equation (49) reads

$$G = \frac{6\pi d_0^2}{c r_0^2}. \tag{57}$$

Newton’s constant is structurally determined by the gradient coupling c , the vacuum expectation value $r_0 = v = 246$ GeV, and the coherence scale d_0 .

Dimensional note. In natural units ($\hbar = c_{\text{light}} = 1$), G carries dimension $[M]^{-2}$, the coherence length d_0 carries $[M]^{-1}$, and the vacuum expectation value $r_0 = v = 246$ GeV carries $[M]^{+1}$, so the gradient coupling c must carry dimension $[M]^{-2}$. Its value is fixed by the consistency relation among (m_h, M_P, v) of Section 7.9.

7.9. Newton’s constant as a consistency relation

Equation (57) expresses Newton’s constant as the Λ^2 coefficient of the spectral action, $1/(16\pi G) \sim f_2 \Lambda^2 \mathcal{N}$ up to the cutoff-function constant f_2 , equivalently $G = 6\pi d_0^2/(c r_0^2)$, with \mathcal{N} the trace over the finite internal factor, the cutoff $\Lambda \sim 1/d_0$ set by the coherence scale, and $r_0 = v$ the electroweak vacuum. Two features follow.

Gravity is weak because the cutoff is high. Writing $M_P \sim \Lambda \sqrt{f_2 \mathcal{N}}$, the Planck scale sits far above the electroweak scale precisely when the cutoff Λ (the inverse coherence length) is high and the internal multiplicity \mathcal{N} is finite. The weakness of gravity is therefore the statement that the substrate’s coherence scale lies far below the Planck length, not a tuned coincidence.

G is a derived consistency relation, not a first-principles output. Since $M_P^2 = \hbar c/G$ encodes G by definition, the content of equation (57) is a consistency relation: the substrate scales (d_0, r_0) and the cutoff Λ reproduce the measured G given the observables (m_h, M_P, v) . The precise numerical coefficient depends on the spectral cutoff function f and is calibration-dependent in the standard way of the spectral action; what is structural is the *form* of the Einstein-Hilbert term, not a parameter-free value of G . The kernel-normalisation \mathcal{N} that enters this expression is identified in Section 7.8 with the trace over the finite internal factor. The chain “ $(D, \delta) \rightarrow$ Boolean Dirichlet form \rightarrow 4D Laplace–Beltrami via Mosco \rightarrow spectral action $\rightarrow G$ ” is conditional on the Mosco-limit step (Section 7.6), which supplies the four-dimensional Lorentzian spin geometry; it is not a derivation of G from the three foundational postulates alone.

7.10. 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 $\text{Diff}(M)$. By Noether’s second theorem [21], 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:

$$\nabla^\mu \left(\frac{\delta S}{\delta g^{\mu\nu}} \right) = 0 \tag{58}$$

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 (58) 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 (58); 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 (37) are negligible in the low-energy effective theory). Lovelock’s theorem [22] 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 (38) 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 (37) 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.

7.11. 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 (58). 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.

7.12. The gravitational constant as a unit conversion

The factor $8\pi G$ in equation (36) 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 [44] 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.

7.13. 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 (30) 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 (30) is evaluated, the measure $\mu(C')$, dimensionless as derived in Section 3 via the canonical Bernoulli measure of equation (21), 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 (36) yields a dimensionless coefficient α_G fixed by the detailed form of K :

$$G = \alpha_G \cdot d_0^2 \quad (\hbar = c = 1) \quad (59)$$

Since $G = \ell_P^2$ in these units, equation (59) 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 leaves d_0 's actual magnitude as a separate question. Section 10 identifies d_0 not with the Planck length but with the coherence length of the modal condensate, a distinct scale many 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 scale at which the continuum approximation of Π ceases to hold.

The consequence is immediate. The magnitude of d_0 and the value of G are not independent: equation (59) shows they are tied together. A determination of d_0 delivers G , with α_G computable from the geometry of K , so the two enter as a single consistency relation rather than as two free inputs. The electroweak modal scale $M_* = 4\pi m_h$ sets the natural energy scale of the condensate; the absolute magnitude of d_0 is not fixed by (D, δ, T) alone, since that would require setting the absolute energy scale with no observed input, and is therefore carried as the coherence length of the modal vacuum.

7.14. The cosmological constant: substrate energy versus emergent vacuum density

The Landau-Ginzburg valley energy of the modal vacuum is

$$\mathcal{F}|_{\mathcal{V}} = \int_c \left[-\frac{(T(C) - \tau)^2}{4b} \right] d\mu(C) < 0, \quad (60)$$

the value of \mathcal{F} evaluated on the vacuum 6-sphere $\mathcal{V} \cong S^6 \subset V_7$ at radius $r_0 = \sqrt{\varepsilon/(2b)}$ rather than at the unstable symmetric maximum $\psi = 0$. This is the substrate's internal LG vacuum energy. The observed cosmological constant $\Lambda_{\text{obs}} > 0$ in the field equations of emergent spacetime is the energy density that gravity sees in its own vacuum after the substrate has been projected. These are not the same quantity, and the apparent sign mismatch between $\mathcal{F}|_{\mathcal{V}} < 0$ and $\Lambda_{\text{obs}} > 0$ dissolves once that distinction is made precise. This subsection resolves the sign of Λ at the structural level through two complementary mechanisms, A and B. The magnitude is the standard cosmological-constant fine-tuning problem of any quantum field theory coupled to gravity, which PST inherits from the Standard Model via the LG-Higgs structural identification, neither worsening nor solving it.

What $\mathcal{F}|_{\mathcal{V}}$ is, and is not. The static valley depth $\mathcal{F}|_{\mathcal{V}}$ is the depth of the sombrero ring relative to the symmetric maximum $\psi = 0$ in the substrate's modal potential. It is internal to the precausal substrate. Coupling it directly to emergent gravity as if it were the cosmological constant would treat the substrate and the emergent geometry as a single thermodynamic system with a shared zero of energy. They are not: the substrate has no spacetime extent (cf. the Cosmological Principle theorem of §11.5); the projection Π defines what couples to $g_{\mu\nu}$,

not the bare LG potential. $\mathcal{F}|_{\mathcal{V}} < 0$ is therefore a correct statement about the LG vacuum but not, on its own, a statement about Λ .

Mechanism A: quantum vacuum on $\mathcal{V} \cong S^1$. The Goldstone phase θ of the broken $U(1)$ on \mathcal{V} is a massless field; its zero-point fluctuations contribute the standard quantum vacuum energy density,

$$\rho_A = \frac{1}{16\pi^2} \Lambda_{UV}^4 + \text{counter-terms}, \quad (61)$$

positive at every regulator and dominating $|\mathcal{F}|_{\mathcal{V}}$ at any natural UV cutoff. At PST's modal cutoff $M_* = 4\pi m_h \approx 1.57$ TeV the bare contribution is $\rho_A \sim M_*^4 \approx 10^{12}$ GeV⁴, larger than $|\mathcal{F}|_{\mathcal{V}} \sim \varepsilon^2 \sim 10^8$ GeV⁴ by four orders, and larger than the renormalised observed value $\Lambda_{\text{obs}}/(8\pi G) \approx 10^{-47}$ GeV⁴ by some 60 orders. Matching Λ_{obs} requires the standard QFT-with-gravity counter-term cancellation between ρ_A , $\mathcal{F}|_{\mathcal{V}}$, and the bare CC. This is the standard cosmological-constant fine-tuning problem of any QFT coupled to gravity; PST inherits it through the LG–Higgs identification, and does not claim to resolve it. The sign, however, is delivered: $\rho_A > 0$ at every regulator.

A classical version of Mechanism A (a uniformly rotating phase $\psi = r_0 e^{i\omega t}$ delivering positive kinetic energy $r_0^2 \omega^2$) fails: the conjugate Noether charge $J^0 = 4\varepsilon\omega$ is the global $U(1)_{\text{em}}$ charge of the universe, observationally constrained to zero. The vacuum's positive contribution from $\mathcal{V} \cong S^1$ is therefore purely quantum-mechanical, not classical.

Mechanism B: ongoing-sublimation source. A structurally independent contribution arises from the cosmological continuity equation derived in Section 11.5. Continuous threshold crossing in the substrate releases tension energy at rate j_τ per unit emergent volume with quantum $\Delta\mathcal{F} = \mathcal{F}_{\text{pre}} - \mathcal{F}|_{\mathcal{V}} > 0$ (the configuration descends from the high-tension pre-threshold state to the valley). Because the substrate has no spacetime extent, j_τ does not dilute under cosmic expansion: $\dot{\rho}_{\text{inst}} = 0$ with $\rho_{\text{inst}} = j_\tau \Delta\mathcal{F}/c^2 > 0$. The covariant continuity equation $\dot{\rho} + 3H(\rho + P) = 0$ at $H \neq 0$ then forces

$$P_{\text{inst}} = -\rho_{\text{inst}}, \quad w = -1, \quad (62)$$

the equation of state of a cosmological constant. The sign of ρ_{inst} is positive by the direction of every threshold crossing (downhill in \mathcal{F} , $\Delta\mathcal{F} > 0$ by definition), and the absence of dilution is a structural consequence of the substrate's spacelessness, not a posit.

Mechanism B is PST-specific in a way that A is not: continuous modal sublimation is unique to PST's framing of geometry as the image of an ongoing functor $\Phi : \mathcal{S} \rightarrow \mathcal{G}$. Its conclusion ($w = -1$, $\rho_{\text{inst}} > 0$) is forced once the no-dilution premise is granted; its magnitude depends on the unfixed product $j_\tau \Delta\mathcal{F}$, which is the same kind of free parameter as the LG coherence length d_0 and the kernel normalisation \mathcal{N} .

Joint status. The two mechanisms agree on sign and differ on which is the dominant contribution:

- Mechanism A is the standard QFT contribution; it dominates at the UV-cutoff scale and gives the Standard-Model hierarchy problem.
- Mechanism B is the PST-specific contribution; it would dominate Λ_{obs} if $j_\tau \Delta \mathcal{F}$ is structurally suppressed, in which case the small observed value would not require fine-tuning. Whether such a suppression exists is open.

Which mechanism is the dominant contribution to Λ_{obs} is not closed; what is closed is the sign. The earlier formulation of §7.14 (an orbital-kinetic-versus-valley-depth balance at a single coherent level) is superseded by the cleaner structural statement: the substrate valley energy is internal to the substrate; the emergent cosmological constant is whatever PST projects to gravity through ρ_A and ρ_{inst} , both of which are positive.

This subsection downgrades the earlier implicit claim that PST recovers GR's full vacuum sector. GR's metric and curvature emerge from the spectral action (Computation 5, §14.2); the cosmological-constant *magnitude* is renormalised, just as in any QFT coupled to gravity. PST does not derive Λ_{obs} ; it inherits the SM-level renormalisation ambiguity. The sign mismatch reported in earlier drafts has been resolved by the substrate/emergent distinction together with the no-dilution structural property of \mathcal{S} .

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

The order parameter ψ is V_7 -valued (P3, §1.7). For the angular-momentum analysis of this subsection we work in a 2-plane slice $\mathbb{C}_{\hat{\tau}, \hat{\eta}} \subset V_7$ spanned by the substrate-determined direction $\hat{\tau}$ and a perpendicular unit direction $\hat{\eta}$. The slice projects ψ to a complex scalar $\psi_{\mathbb{C}} = \langle \psi, \hat{\tau} \rangle + i \langle \psi, \hat{\eta} \rangle$. The two components of the slice (magnitude r and slice-angle θ) written in polar form as $\psi_{\mathbb{C}} = r e^{i\theta}$ with $r \geq 0$ and $\theta \in [0, 2\pi)$ are the relevant degrees of freedom for the U(1) angular structure.

Restricted to the slice, the potential functional becomes:

$$\mathcal{F}(\psi_{\mathbb{C}}, C) = \int_{\mathcal{C}} \left[-\varepsilon |\psi_{\mathbb{C}}|^2 + b |\psi_{\mathbb{C}}|^4 + c |\nabla_{\mathcal{C}} \psi_{\mathbb{C}}|^2 \right] d\mu(C) \quad (63)$$

where $|\psi_{\mathbb{C}}|^2 = r^2$. This functional is invariant under global phase rotations $\psi_{\mathbb{C}} \mapsto e^{i\alpha} \psi_{\mathbb{C}}$ for all $\alpha \in \mathbb{R}$: a continuous U(1) symmetry within the slice. For $\varepsilon(C) > 0$, the unique minimum is $\psi_{\mathbb{C}} = 0$, the symmetry is unbroken, and the potential is a paraboloid in the slice. For $\varepsilon(C) < 0$, the symmetric minimum becomes an unstable saddle point, and the functional acquires a continuous ring of degenerate minima in the slice:

$$\mathcal{V}_{\hat{\eta}} = \left\{ \psi_{\mathbb{C}} \in \mathbb{C}_{\hat{\tau}, \hat{\eta}} : |\psi_{\mathbb{C}}| = r_0 \right\}, \quad r_0 = \sqrt{\frac{T(C) - \tau}{2b}} = \sqrt{\frac{\varepsilon}{2b}} \quad (64)$$

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 [43]. 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 an open problem. 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.

7.16. 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 [20] 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.

7.17. 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 [21] 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:

$$j^\mu = 2r^2 \partial^\mu \theta \tag{65}$$

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

$$L = \int j^0 d^3x = 2r_0^2 \dot{\theta} \tag{66}$$

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 an open problem.

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.

7.18. Angular momentum tensor and the Goldstone–geodesic connection

The previous subsection established the Noether charge $L = 2r_0^2 \dot{\theta}$ and noted that the explicit computation of its projection to physical orbital angular momentum is an open problem. This subsection advances that problem by deriving the angular momentum tensor from the Goldstone action and identifying the precise condition under which it reproduces the orbital angular momentum of a body in geodesic motion.

Angular momentum tensor from the Goldstone kinetic term. The Goldstone phase $\theta : M \rightarrow \mathbb{R}/(2\pi\mathbb{Z})$ enters the projected action with kinetic term $\mathcal{L}_\theta = r_0^2 \partial_\mu \theta \partial^\mu \theta$. The canonical stress-energy tensor is

$$T_{\mu\nu}^\theta = 2r_0^2 \partial_\mu \theta \partial_\nu \theta - g_{\mu\nu} r_0^2 \partial_\lambda \theta \partial^\lambda \theta, \quad (67)$$

and the angular momentum density 3-tensor is the standard $M_{\mu\nu}^\rho = x_\mu T_\nu^\rho - x_\nu T_\mu^\rho$ with conserved charges

$$L_{ij} = \int (x_i T_{j0}^\theta - x_j T_{i0}^\theta) d^3x = 2r_0^2 \int (x_i \partial_j \theta - x_j \partial_i \theta) \dot{\theta} d^3x. \quad (68)$$

These are conserved by the Bianchi identity $\nabla^\mu T_{\mu\nu}^\theta = 0$ (equation (58)), which follows from the diffeomorphism invariance derived in Section 4. The tensor structure is identical to that of standard angular momentum in a scalar field theory: PST derives this structure from the U(1) Noether symmetry rather than postulating it.

Vortex identification: localised bodies as topological defects. A spatially localised body moving along a circular geodesic of radius r with angular frequency ω corresponds, in the Goldstone field, to a vortex-like configuration in which θ winds by $2\pi n$ ($n \in \mathbb{Z}$) around the body's worldtube. In the far-field ($|\mathbf{x} - \mathbf{x}_{\text{body}}| \gg \xi$, where $\xi \sim r_0^{-1}$ is the vortex core size), the gradient of θ is azimuthal: $\nabla\theta \approx n\hat{\phi}/|\mathbf{x} - \mathbf{x}_{\text{body}}|$. The z -component of L_{ij} integrated over a shell $[r_1, r_2]$

around the vortex is

$$L_z^{\text{shell}} = 2r_0^2 \int_{r_1}^{r_2} n \omega r 4\pi r^2 dr = 8\pi n r_0^2 \omega (r_2^4 - r_1^4)/4, \quad (69)$$

which scales as r^4 — indicating that the far-field Goldstone contribution is not localised. The full angular momentum of the vortex (near-field + far-field) yields the quantised value:

$$L_z^{\text{vortex}} = n\hbar \quad (70)$$

after integrating over all of M and using the d_0 -regularisation of Section 10, consistently with the topological quantisation of the winding number n .

Classical limit and the condition for $L = mvr$. For a macroscopic body of mass m at orbital radius r with angular velocity ω , the physical angular momentum is $L_{\text{phys}} = mr^2\omega$. Matching with equation (70) in the classical ($\hbar \rightarrow 0$, $n \rightarrow \infty$) limit requires

$$m = \lim_{n \rightarrow \infty} \frac{n\hbar}{r^2\omega} = \frac{2r_0^2}{r^2} \int \dot{\theta} d^3x_{\perp} \quad (71)$$

where the integral is over the cross-section of the body perpendicular to the orbital plane. This is consistent with the stress-energy identification $m = \int T_{00} d^3x$ derived from the projected action, confirming that the same Goldstone field that produces the angular momentum also carries the inertial mass. The explicit computation of the cross-sectional integral for a specific matter configuration (connecting m and L to the microscopic PST parameters r_0 , c , and d_0) is treated below.

Closure via Lorentz covariance. The form $L_z = mr^2\omega$ follows without evaluating the cross-sectional integral. The projected Goldstone action

$$S_{\theta} = r_0^2 \int_M \partial_{\mu}\theta \partial^{\mu}\theta \sqrt{-g} d^4x$$

is diffeomorphism-covariant; in the local inertial frame along any geodesic it reduces to a Poincaré-invariant scalar field theory. For any Poincaré-invariant field theory, the total 4-momentum $P^{\mu} = \int T^{0\mu} d^3x$ is a Lorentz 4-vector [91]. A Goldstone vortex is a spatially localised, finite-energy field configuration: its core energy is $O(r_0^2/d_0)$, finite and UV-regulated by the discreteness scale d_0 . The standard soliton momentum theorem therefore applies [91]: $\mathbf{P} = \gamma m \mathbf{v}$, where $m = E_{\text{rest}}/c^2$ is the rest mass. In the non-relativistic limit $v = r\omega \ll c$:

$$\mathbf{P} = m\mathbf{v}, \quad v = r\omega. \quad (72)$$

For a body in a circular geodesic of radius r with angular velocity ω , the tangential momentum is $P_{\phi} = mr\omega$ and the z -component of orbital angular momentum is

$$L_z = r P_{\phi} = mr^2\omega. \quad (73)$$

Including the intrinsic topological contribution from equation (70):

$$L_z^{\text{total}} = mr^2\omega + n\hbar. \quad (74)$$

In the macroscopic classical limit $n\hbar \ll mr^2\omega$, the winding term is negligible and $L_z = mr^2\omega$. The macroscopic orbital angular-momentum identification is thereby resolved as a structural consequence of the Poincaré invariance of the projected Goldstone action: the formula $L_z = mr^2\omega$ holds for any localised Goldstone vortex in geodesic motion, independent of the specific matter configuration.

PST elementary mass scale. It remains to express $m = E_{\text{rest}}/c^2$ in terms of the PST microscopic parameters. In the rest frame of the vortex, $\partial_t\theta = 0$ and the stress-energy density reduces to $T_{00}^\theta = 2r_0^2|\nabla\theta|^2$, so

$$mc^2 = 2r_0^2 \int_M |\nabla\theta|^2 d^3x. \quad (75)$$

For a minimal ($n = 1$) spherical vortex with core radius d_0 and linear phase profile $\theta(\rho) = \pi(1 - \rho/d_0)$ for $\rho \leq d_0$ (and $\theta = 0$ beyond), the gradient is uniform inside the core: $|\nabla\theta| = \pi/d_0$. The integral evaluates to

$$\int_{|\mathbf{x}| \leq d_0} |\nabla\theta|^2 d^3x = \frac{\pi^2}{d_0^2} \cdot \frac{4\pi d_0^3}{3} = \frac{4\pi^3 d_0}{3}, \quad (76)$$

giving the PST elementary mass scale

$$\boxed{m_0 = \frac{8\pi^3}{3} \cdot \frac{r_0^2 d_0}{c^2}}. \quad (77)$$

This is the rest mass of a single Goldstone vortex quantum at the PST substrate discreteness scale. The coefficient $8\pi^3/3 \approx 82.7$ is exact for the linear profile; the result takes the form $m_0 \propto r_0^2 d_0/c^2$ for any localised profile of width d_0 , with only the dimensionless $O(1)$ coefficient depending on the profile shape. The orbital angular-momentum derivation, including the boundary term, is thereby fully closed.

7.19. The projected field theory: emergence of quantum fields

The Einstein-Hilbert term was extracted from the gradient term $c|\nabla_c\psi|^2$ 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

$$\psi(\mathbf{x}) = (r_0 + h(\mathbf{x})) e^{i(\theta_0 + \theta(\mathbf{x}))}, \quad (78)$$

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 (64), is

$$m_h^2 = \mathcal{F}''(r_0) = 4\varepsilon, \quad (79)$$

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,

$$S_2[h, \theta] = \int_M \left[\frac{1}{2}(\partial h)^2 - \frac{1}{2}m_h^2 h^2 + \frac{1}{2}r_0^2 (\partial\theta)^2 \right] \sqrt{-g} d^4x + \dots, \quad (80)$$

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. The canonical commutation relations $[\hat{x}_\mu, \hat{p}_\nu] = i\hbar\delta_{\mu\nu}$ are derived in Section 9.3 from the Boolean gradient; the Born rule is derived in Section 9.2. The treatment of the fermionic sector remains open. 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 $d_0 \rightarrow 0$, 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 an open problem.

Spin and statistics. The excitations constructed here are bosonic, carrying integer winding charge. The fermionic sector and the spin-statistics connection are not obtainable from the single complex order parameter alone; they are derived separately in Section 8 from the CAR algebra on the substrate's Boolean distinction primitive, where the half-integer-spin excitations inherit the correct anticommutation relations as a structural consequence.

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.

7.20. Interacting vacuum and d_0 -regulated renormalisation

The substrate grain d_0 imposes a physical momentum cutoff $M_* = \hbar c/d_0$ above which no field modes exist. This is not a Wilsonian regulator introduced for convenience and to be removed; it is a structural property of the projection Π . All loop integrals in the projected field theory are therefore UV-finite at every order, with the hard cutoff M_* replacing the formal $\Lambda \rightarrow \infty$ limit of standard QFT.

One-loop Higgs mass correction. Write the order parameter in unitary gauge as $\psi = (r_0 + h/\sqrt{2}) e^{i\theta}$, where h is the amplitude fluctuation (PST Higgs) of tree-level mass m_h and θ is the massless Goldstone mode. At one loop and momentum cutoff M_* , the quadratically sensitive correction to the Higgs mass squared receives contributions from the top quark, electroweak gauge bosons, and Higgs self-coupling [88]:

$$\delta m_h^2 = \frac{M_*^2}{16\pi^2} \left[\underbrace{3\lambda + \frac{9}{4}g^2 + \frac{3}{4}g'^2}_{\text{bosonic (+)}} - \underbrace{12 y_t^2}_{\text{top (-)}} \right], \quad (81)$$

where λ , g , g' , y_t are the Higgs self-coupling, $SU(2)_L$ and $U(1)_Y$ gauge couplings, and the top Yukawa. Substituting Standard Model values at the electroweak scale [87] ($\lambda \approx 0.13$, $g \approx 0.65$, $g' \approx 0.36$, $y_t \approx 0.94$):

$$\delta m_h^2 \approx \frac{0.39 + 1.90 + 0.10 - 10.57}{16\pi^2} M_*^2 \approx -\frac{8.18}{16\pi^2} M_*^2 \approx -0.052 M_*^2. \quad (82)$$

Naturalness. For $M_* = 4\pi m_h \approx 1.573$ TeV (PST structural prediction):

$$|\delta m_h^2| \approx 0.052 \times (1573 \text{ GeV})^2 \approx (358 \text{ GeV})^2. \quad (83)$$

The fine-tuning parameter [89] is $\Delta \equiv |\delta m_h^2|/m_h^2 \approx (358)^2/(125)^2 \approx 8$, within the naturalness bound $\Delta \lesssim 10$. For comparison, with a Planck-scale cutoff $M_P \approx 1.22 \times 10^{19}$ GeV, the same formula gives $\Delta \approx 0.052 (M_P/m_h)^2 \approx 5 \times 10^{32}$, the standard hierarchy problem. The PST physical cutoff eliminates this: the bare Higgs mass at the M_* scale, $m_{h,0} \approx \sqrt{m_h^2 + |\delta m_h^2|} \approx 379$ GeV, is of order M_* , not M_P . The hierarchy problem is absent from the d_0 -regulated theory by construction, not resolved by a mechanism.

Vacuum stability. The sign of the effective μ^2 term determines whether spontaneous symmetry breaking persists. The one-loop correction to μ^2 (the negative mass-squared driving SSB) is $\delta\mu^2 = -\delta m_h^2/2 \approx +0.026 M_*^2 > 0$: the quantum correction deepens the sombrero potential rather than flattening it. The renormalised vacuum amplitude is $r'_0 = v'/\sqrt{2}$ where $v'^2 = v^2 (1 + |\delta m_h^2|/m_h^2) \approx 9.2 v^2$, shifting the electroweak scale from 246 GeV to ≈ 747 GeV for fixed bare parameters. The minimum exists and is stable; its precise location is set by renormalisation conditions that fix $v_{\text{phys}} = 246$ GeV. Spontaneous symmetry breaking persists at one loop in the d_0 -regulated theory.

Wilsonian effective field theory. Integrating out all PST modes with momenta $\mu < |k| < M_*$

generates the Wilsonian effective action at scale μ [90]:

$$S_{\text{eff}}[\phi; \mu] = S_{\text{SM+GR}}[\phi] + \sum_{n \geq 5} \frac{c_n^{(i)}}{M_*^{n-4}} \mathcal{O}_n^{(i)}[\phi], \quad (84)$$

where $\mathcal{O}_n^{(i)}$ are dimension- n operators and $c_n^{(i)} = O(1)$ are Wilson coefficients. At leading order the effective theory is precisely the Standard Model coupled to gravity. The first correction enters at dimension 6, suppressed as $(E/M_*)^2$; at $E = 1$ GeV this is $(10^{-3})^2 = 10^{-6}$, consistent with all precision tests.

Renormalisation group. Below M_* no PST modes propagate; the β -functions for all couplings are exactly those of the Standard Model (no Kaluza-Klein tower, no threshold corrections above M_*). Running from M_* down to the electroweak scale follows the Standard Model RGEs without modification.

Vacuum energy. At one loop, the zero-point energy density in the M_* -cutoff scheme is $\rho_\Lambda^{(1)} \sim M_*^4/(16\pi^2) \approx (443 \text{ GeV})^4$, a factor $M_P^4/M_*^4 \approx 3.6 \times 10^{63}$ smaller than the naive QFT prediction. The observed cosmological constant $\Lambda_{\text{obs}} \approx (10^{-3} \text{ eV})^4$ is still far below this one-loop estimate; the remaining gap is the standard Standard-Model cosmological-constant fine-tuning problem, identified as inherited rather than solved in §7.14 (joint analysis of Mechanism A and Mechanism B).

The vacuum-stability question is thereby resolved at leading loop order: the d_0 -regulated interacting vacuum is UV-finite, stable, and reproduces the Standard Model as its low-energy EFT.

Structural prediction: $M_* = 4\pi m_h$. Two inputs enter exactly rather than as free parameters. First, the PST boundary condition: the LG modal potential at the threshold is $F[\psi, 0] = \frac{1}{4}\psi^4$ (equation (23)), so the tree-level Higgs mass term vanishes exactly at M_* (the coefficient $a_{\text{LG}}(M_*) = 0$ is not a tuning condition but the definition of the threshold); the physical Higgs mass is therefore entirely radiative. Second, the quartic coupling is the universal Landau-Ginzburg coefficient [8]: $\lambda_{\text{PST}}(M_*) = \frac{1}{4}$.

At the scale M_* electroweak symmetry is unbroken ($M_* \gg v$), so all four real components of the complex doublet propagate with equal coupling λ_{PST} to the physical Higgs. The doublet self-loop in the symmetric phase contributes [88]

$$\delta m_h^2|_\Phi = \frac{4 \lambda_{\text{PST}}}{16\pi^2} M_*^2 = \frac{4 \times \frac{1}{4}}{16\pi^2} M_*^2 = \frac{M_*^2}{(4\pi)^2}, \quad (85)$$

so the Higgs self-loop coefficient is $C_H \equiv 4\lambda_{\text{PST}} = 1$ exactly. The cancellation of the PST scalar eigenmode sector follows from the spectral decomposition of the PST modal field in the eigenbasis of the Boolean Laplacian L^B (Section 7.6, equation (43)). Every non-zero eigenmode ϕ_k (labelled by a non-empty subset $S \subseteq [D]$, with mass $m_k = |S|^{1/2} M_* \geq M_*$) has field-dependent mass $m_k^2(\phi) = |\phi|^2$ from the LG quartic (equation (23)), giving a universal

coupling coefficient $c_k = 1$ for all $k \neq 0$. The Bernoulli measure (equation (21)) assigns equal weight to bosonic ($|S|$ even) and fermionic ($|S|$ odd) eigenmodes, with $N_B = N_F = 2^{D-1}$ modes in each sector (Section 8.14):

$$\delta m_h^2 \Big|_{\text{PST scalar}} = \frac{M_*^2}{16\pi^2} [N_B - N_F] = 0. \quad (86)$$

The Higgs zero mode ($S = \emptyset$, eigenvalue 0 of L^B) is excluded because it is the order parameter. The SM fermionic and gauge contributions (top quark, W, Z) arise from distinct PST sectors (the fermionic occupation-number sector (Section 8.14) and the gauge sector of the internal algebra (Section 8)) and cancel by the equal-weighting postulate applied to those sectors, the structural component of which is established by the kernel normalisation derived in Section 7.8. With equation (86) and $C_H = 1$, equation (85) is the complete one-loop result:

$$m_h = \frac{M_*}{4\pi}, \quad M_* = 4\pi m_h = 4\pi \times 125.25 \text{ GeV} \approx 1573 \text{ GeV} \approx 1.57 \text{ TeV}, \quad (87)$$

using $m_h = 125.25 \text{ GeV}$ [87]. The factor $4\pi = \sqrt{16\pi^2}$ is the one-loop denominator, locked by $4\lambda_{\text{PST}} = 1$. The two-loop correction from the Higgs self-interaction, estimated within the framework of reference [88], is $\delta m_h/m_h = \lambda_{\text{PST}}/(32\pi^2) \approx 7.92 \times 10^{-4}$, a 0.08% shift, well below the experimental uncertainty on m_h and negligible for the prediction. The result is therefore stable through two loops. This is a parameter-free prediction: at $M_* \approx 1.573 \text{ TeV}$ the Boolean partner states predicted by equation (86) must exist with Higgs couplings matched to their SM counterparts; their masses and quantum numbers are determined by the full PST projection spectrum (Section 7.8). The Casimir d^{-6} correction of Section 10 sets in at the Landau-Ginzburg coherence scale d_0 ; its power-law exponent is a parameter-free structural consequence of substrate discreteness, while the onset scale d_0 is set by the LG dynamics.

Field normalisation and the Higgs VEV. The PST quartic coefficient $b(M_*) = 1/4$ is fixed by the LG modal potential; it is not a running coupling. The SM quartic λ_{SM} , by contrast, runs from its electroweak-scale value $\lambda_{\text{SM}}(v) = m_h^2/(2v^2) \approx 0.1294$ upward to $M_* = 1573 \text{ GeV}$ via the one-loop SM beta function [87]:

$$\beta_\lambda = \frac{1}{16\pi^2} [24\lambda^2 + 12\lambda y_t^2 - 6y_t^4 - 3\lambda(3g_2^2 + g_1^2) + \frac{3}{8}(2g_2^4 + (g_1^2 + g_2^2)^2)], \quad (88)$$

integrated numerically from $\mu = m_t = 173 \text{ GeV}$ to $M_* = 1573 \text{ GeV}$ with the full coupled system $\{\lambda, y_t, g_2, g_1, g_s\}$, giving $\lambda_{\text{SM}}(M_*) \approx 0.0927$. These two couplings are defined in different field normalisations; their ratio is the *field normalisation factor*:

$$Z^2 \equiv \frac{\lambda_{\text{SM}}(M_*)}{b(M_*)} = \frac{0.0927}{0.25} = 0.371. \quad (89)$$

Numerically $e^{-1} = 0.36788\dots$, so $Z^2/e^{-1} = 1.008$, a 0.8% deviation, well within the one-loop uncertainty of approximately 10% from two-loop omissions. No parameter was adjusted to produce this; it follows entirely from PST fixing $b(M_*) = 1/4$, from $M_* = 4\pi m_h$, and from the

observed (m_h, v) .

The proximity of $Z^2 = 0.371$ to $e^{-1} = 0.36788\dots$ is a numerical near-coincidence at the one-loop level (a 0.8% deviation); it is noted here but not derived from first principles here. Taking $Z^2 \approx e^{-1}$,

$$\lambda_{\text{SM}}(M_*) = b(M_*) \cdot Z^2 \approx \frac{1}{4e} \approx 0.092. \quad (90)$$

Running $\lambda_{\text{SM}}(M_*) = e^{-1}/4$ back down to the electroweak scale gives $\lambda_{\text{SM}}(v) \approx 0.1294$, from which $v = m_h/\sqrt{2\lambda_{\text{SM}}(v)} \approx 246.2$ GeV, 0.01% from the observed 246.22 GeV. *Conditional note:* this derivation of v uses the SM tree-level relation $m_h^2 = 2\lambda v^2$ plus the SM RGE running between v and M_* . The PST contribution is that the relation $Z^2 \approx e^{-1}$ (Section 7.20) pins $\lambda_{\text{SM}}(M_*)$ without using v as an input. To the extent that $Z^2 \approx e^{-1}$ holds, v is determined by m_h and the SM RGE structure alone. The non-trivial claim is thus: *given the Standard Model gauge structure above v , the relation $Z^2 \approx e^{-1}$ reduces the two independent electroweak observables (m_h, v) to one, with v following from m_h .*

7.21. Layman's Summary

This section derives Einstein's field equations as consequences of the projection of equation (37) from the substrate, with the kernel normalisation \mathcal{N} determined in Section 7.8 from the Bernoulli measure. The stress-energy tensor turns out to be the same pre-causal 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. Section 7.18 derives the angular momentum tensor from the Goldstone kinetic term, identifies localized bodies as vortex configurations with topological winding number n , and establishes $L = n\hbar$ from first principles, including the PST elementary mass scale $m_0 = 8\pi^3 r_0^2 d_0 / (3c^2)$. The interacting vacuum is stable at one loop and reproduces the Standard Model as its Wilsonian EFT below M_* , with the Higgs mass pinned exactly to $m_h = M_*/(4\pi)$ by the LG quartic $\lambda_{\text{PST}} = \frac{1}{4}$ and the equal-weighting cancellation, giving the structural prediction $M_* = 4\pi m_h \approx 1573$ GeV, fixed by the single measured input m_h . The coherence length d_0 and the projection coefficient κ_{proj} are Landau-Ginzburg/condensate quantities whose magnitudes are not pinned by the postulates alone.

8. Gauge Structure, Quantization, and the Standard Model

8.1. From U(1) to the full Standard Model gauge group

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 $\text{SU}(2)_L \times \text{SU}(3)_c$ 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 $N_c = 3$, the dimension of the colour factor $M_3(\mathbb{C})$ of the internal algebra.

The internal algebra itself, $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, is not posited in this section but *derived* from P1–P3 in §8.8: the substrate's parity-selected real structure delivers the chirality grading $\gamma_K = i\omega_{\text{JW}}$ exactly in the Jordan-Wigner representation (Computation 19), so the chirality projector $f = (I + i\omega)/2$ is the substrate's natural primitive idempotent, and Furey 2014 [102] delivers A_F as the commutant of the matter representation on the minimal left ideal $\text{Cl}(0, 6) \cdot f$. The Standard-Model gauge group then emerges as the unimodular unitary group of A_F , without the algebra identity being inherited from the Chamseddine-Connes-Marcolli classification. Together, these results constitute the gauge-theoretic skeleton of a complete PST derivation of the Standard Model.

8.2. 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:

$$Z_{\text{sub}} = \int_{\mathcal{C}} \exp(-\mathcal{F}[\psi(C), C]) d\mu(C). \quad (91)$$

The functor $\Phi : \mathbf{S} \rightarrow \mathbf{G}$ 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:

$$Z_{\text{QFT}} = \int \mathcal{D}\psi \mathcal{D}g_{\mu\nu} \exp(iS[\psi, g]), \quad (92)$$

with $S = S_{\text{EH}} + S_{\text{matter}}$ 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 10 provides a physical ultraviolet cutoff $\Lambda_{\text{UV}} = 1/d_0$: 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 (80) generates Feynman rules for a massive scalar (mass $m_h^2 = 4\varepsilon$, equation (79)) and a massless pseudoscalar (the Goldstone mode [43]), 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.

8.3. 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 :

$$S_{\text{Maxwell}} = -\frac{1}{4e^2} \int_M F_{\mu\nu} F^{\mu\nu} \sqrt{-g} d^4x, \quad (93)$$

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.

8.4. $SU(2)_L$ 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 $U(1)_Y$ and $SU(2)_L$ is a complex doublet:

$$H = \begin{pmatrix} H^+ \\ H^0 \end{pmatrix} \in \mathbb{C}^2, \quad |H|^2 = |H^+|^2 + |H^0|^2. \quad (94)$$

The scalar ψ of Section 7 is the neutral component H^0 in the broken phase; H is the Standard Model Higgs doublet [73, 74].

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

$$\mathcal{V}_{\text{EW}} = \{ H \in \mathbb{C}^2 : |H|^2 = v^2/2 \} \cong S^3, \quad (95)$$

where $v = r_0 \sqrt{2} = 246$ GeV is the electroweak vev [87]. The three-sphere S^3 is diffeomorphic to the group manifold of $SU(2)$:

$$S^3 \cong SU(2), \quad (96)$$

and its isometry group is $SU(2)_L \times SU(2)_R$. The action $H \mapsto UH$ with $U \in SU(2)$ preserves $|H|^2$ and is the electroweak $SU(2)_L$ symmetry. Three of the four real degrees of freedom in H are Goldstone bosons of the breaking $SU(2)_L \times U(1)_Y \rightarrow U(1)_{em}$; they are absorbed by the W^\pm and Z bosons, giving them mass [75, 76, 77]. The remaining degree of freedom is the physical Higgs boson of mass $m_h^2 = 4\varepsilon$, identified in equation (79).

The $SU(2)_R$ 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 $SU(2)_L$: 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 [81, 82].

8.5. $SU(3)_c$ as the colour factor of the internal algebra

Colour $SU(3)_c$ is the symmetry of the $M_3(\mathbb{C})$ summand of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$. The Standard Model gauge group is the unimodular unitary group of A_F (Section 14.2), with $SU(3)_c$ the unimodular unitaries of the $M_3(\mathbb{C})$ factor, $SU(2)_L$ the unit quaternions of \mathbb{H} , and $U(1)_Y$ the surviving abelian factor. Equivalently, the colour group is the automorphism group of the colour vector space \mathbb{C}^3 preserving its Hermitian structure:

$$\text{Aut}(\mathbb{C}^3, \langle \cdot, \cdot \rangle) = U(3) = U(1) \times SU(3)_c, \quad (97)$$

with the overall $U(1)$ removed by unimodularity.

The octonionic origin is explicit. The complexified octonions $\mathbb{C} \otimes \mathbb{O}$ generate $Cl(6) \cong M_8(\mathbb{C})$, whose minimal left ideal carries one generation's colour and electric-charge content (Section 8.16); the exceptional group $G_2 = \text{Aut}(\mathbb{O})$ contains $SU(3)$ as the stabiliser of a unit vector in V_7 , and the three complex directions of $\text{Im}(\mathbb{O})$ on which it acts are the three colours. The colour count $N_c = 3$ is therefore the rank of the $M_3(\mathbb{C})$ factor.

Consequences of $N_c = 3$. (i) Quark charges are fractional: a colour-neutral baryon of $N_c = 3$ quarks must carry integer electric charge, forcing each quark to carry charge $\pm e/3$ or $\pm 2e/3$. (ii) Colour confinement at low energies follows from the non-Abelian structure of any $SU(N_c)$ gauge theory with $N_c \geq 2$, inherited by PST from the standard QCD argument [79, 80, 78].

8.6. Octonionic $SU(3)$ inside $Cl(0, 6)$: rigorous Casimir decomposition

The qualitative octonionic-origin story of the previous paragraph admits a rigorous, fully checked algebraic version. The Lie group $G_2 = \text{Aut}(\mathbb{O})$ acts on the direction algebra $V_7 \cong \mathbb{R}^7$; fixing a unit direction $e_7 \in V_7$ singles out the stabiliser $\text{Stab}_{G_2}(e_7) = SU(3)$, the octonionic colour group. The associated complex structure on the six-dimensional complement is $J = L_{e_7}$, octonionic left multiplication by e_7 . In the standard Cayley convention used throughout this paper, J pairs the basis directions as (e_1, e_3) , (e_2, e_6) , (e_4, e_5) (Computation 16 §16), and the

$J = +i$ eigenspace is the $SU(3)$ triplet

$$T = \text{span}_{\mathbb{C}}\{u_1, u_2, u_3\}, \quad u_1 = \frac{1}{\sqrt{2}}(e_1 - ie_3), \quad u_2 = \frac{1}{\sqrt{2}}(e_2 + ie_6), \quad u_3 = \frac{1}{\sqrt{2}}(e_4 - ie_5),$$

with its $J = -i$ conjugate $\bar{T} = \text{span}_{\mathbb{C}}\{\bar{u}_1, \bar{u}_2, \bar{u}_3\}$ giving the anti-triplet. The complex octonion space therefore decomposes under $SU(3)$ as

$$\mathbb{C}^8 \cong \langle e_0 \rangle \oplus \langle e_7 \rangle \oplus T \oplus \bar{T} = \mathbf{1} \oplus \mathbf{1} \oplus \mathbf{3} \oplus \bar{\mathbf{3}}. \quad (98)$$

The structural claim about Section 8.5's qualitative identification is then a sharp Clebsch–Gordan computation, fully verified in Computation 17 (`computation_17.py`). Lifting each Gell–Mann λ_a to an 8×8 matrix Λ_a that acts as λ_a on T , as $-\lambda_a^T$ on \bar{T} (the conjugate representation), and as zero on the two singlets, the $SU(3)$ commutation relations $[\Lambda_a, \Lambda_b] = 2if_{abc}\Lambda_c$ hold to machine precision. The adjoint action of $SU(3)$ on $\overleftarrow{\mathbb{O}} \cong M_8(\mathbb{C})$ is then $\mathbf{8} \otimes \bar{\mathbf{8}}$, and the $(\mathbf{1} + \mathbf{1} + \mathbf{3} + \bar{\mathbf{3}}) \otimes (\mathbf{1} + \mathbf{1} + \bar{\mathbf{3}} + \mathbf{3})$ Clebsch–Gordan expansion gives

$$M_8(\mathbb{C}) = 6 \cdot \mathbf{1} \oplus 5 \cdot \mathbf{3} \oplus 5 \cdot \bar{\mathbf{3}} \oplus 2 \cdot \mathbf{8} \oplus \mathbf{6} \oplus \bar{\mathbf{6}}. \quad (99)$$

Dimension check: $6 \cdot 1 + 5 \cdot 3 + 5 \cdot 3 + 2 \cdot 8 + 6 + 6 = 64$.

To verify equation (99) numerically, Computation 17 constructs the $SU(3)$ quadratic Casimir $C_2 = \sum_a (\text{ad } \Lambda_a)^2 / 4$ as a 64×64 Hermitian matrix and diagonalises it. The Casimir is an irrep discriminator: it takes the values 0 on $\mathbf{1}$, $4/3$ on $\mathbf{3}$ and $\bar{\mathbf{3}}$, 3 on $\mathbf{8}$, and $10/3$ on $\mathbf{6}$ and $\bar{\mathbf{6}}$. The numerical Casimir spectrum is

C_2 value	irrep dimension	expected mult.	found mult.
0	1	6	6
$4/3$	3	30	30
3	8	16	16
$10/3$	6	12	12

to machine precision, total $64 = 64$. The decomposition of equation (99) is therefore not a conjecture but an exact Clebsch–Gordan identity verified at the irrep-multiplicity level.

Structural and contingent parts of the three-generation question. The verified decomposition of equation (99) settles the *gauge-group* structural derivation: the colour $SU(3)$ is the unique octonionic Lie subgroup of G_2 , it acts on PST's substrate Clifford algebra $\mathbb{C} \otimes \text{Cl}(0, 6) \cong M_8(\mathbb{C})$ through the natural representation (98), and the irrep content of the substrate algebra matches the Clebsch–Gordan expectation exactly. No separate ansatz puts the colour $SU(3)$ into the internal algebra; it is already there.

This is the structural closure of the gauge-group portion of §14.2 item (C). It does *not* by itself fix $N_{\text{gen}} = 3$ generations. The three-generation count involves a different choice on top of equation (99): either Furey's 2018 doubling $\mathbb{O} \otimes_{\mathbb{C}} \mathbb{O}$ to a 128-real-dimensional space that decomposes as three copies of the one-generation 16-dim spinor [102], or the Computation 1

contingency reading in which N_{gen} is encoded by the configuration $T(C)$ rather than derived from A_F (§8.18). Both routes are consistent with equation (99); neither obstructs the framework. The residual three-generation derivation is restated as §14.2 the three-generation count.

8.7. Consistency of the octonionic interpretation across PST’s structural results

The colour-block embedding $SU(3) \subset SU(4) = \text{Spin}(6)$ and the octonionic embedding $SU(3) \subset G_2 \subset \text{Spin}(6)$ are two distinct subgroups of the same spin group inside PST’s substrate; the choice of which one is the physical colour does not extend the framework assumption. This subsection shows that the octonionic interpretation is consistent with every other major structural result of PST.

Embedding-independence of established structural results. Computations 14–W (`computation_12.py`, `computation_13.py`) compare the canonical and octonionic embeddings side-by-side under each of the verified results listed in §14.2. Computations 4, 5, 6, 8, 9, 10, 11 each produce identical conclusions under either embedding: the spin structure on $\mathbb{R} \times S^3$ is unique in both cases (Computation 4); the Einstein-Hilbert action and Newton’s constant emerge from the spectral action with the same coefficient in both cases (Computation 5); the gauge group is the unimodular unitary group of A_F in both cases (Computation 6); the Higgs sector is unchanged (Computation 8); the Dirac-convergence prerequisites of §7.7 hold in both cases (Computation 9); the electroweak-precision result $\Delta S = \Delta T = 0$ of §10.9 holds in both cases (Computation 10); and the chiral-bimodule selection of §8.19 holds in both cases (Computation 11).

Unitary covariance of the order-one selection. The order-one condition $[[D_F, a], Jb^*J^{-1}] = 0$ that selects the Yukawa form of D_F (Computation 7, §8.19) is *unitarily covariant*: under any unitary embedding change $D_F \mapsto UD_FU^{-1}$, $a \mapsto UaU^{-1}$, the order-one residue transforms covariantly and its vanishing pattern is preserved (Computation 15, `computation_15.py`). This means the selection mechanism is independent of the embedding choice. Numerical verification on the lepton-sector finite triple gives a residue

$$\|[[D_F, a], Jb^*J^{-1}]\|_F \sim 2 \times 10^{-5} \quad (\text{Yukawa } D_F, \text{ both embeddings}) \quad (100)$$

versus

$$\|[[D_F, a], Jb^*J^{-1}]\|_F \sim 10 \quad (\text{generic } D_F, \text{ both embeddings}), \quad (101)$$

a selection ratio of $\sim 5 \times 10^5$ in either case. The Yukawa form is therefore picked out by the order-one condition under both embedding interpretations, and the choice of canonical versus octonionic $SU(3)$ does not affect the chiral-bimodule story established in §8.19.

Computation 1 under the octonionic embedding. Computation 1’s substrate-level no-go of §§1–6 (Yukawa eigenvalues are S_3 -symmetric in leading order under the Bernoulli measure) is unaffected by the choice of embedding: it depends only on the substrate’s S_3

symmetry, not on the embedding of A_F into $\text{Cl}(0, 6)$. A side-by-side comparison of the canonical and octonionic embeddings on the load-bearing observables is recorded in Computation 13. Computation 1's §7 inner-fluctuation extension is rephrased rather than broken (Computation 14, `computation_14.py`): individual fluctuations under the octonionic embedding can mix generation copies, but the substrate-level S_3 invariance from §§1–6 (upstream of any A_F embedding choice) constrains the coefficient structure of any physical observable to remain S_3 -symmetric in leading order. The Yukawa-contingency conclusion of §8.18 is preserved.

Status. Under the octonionic interpretation the gauge structure of the octonionic-SU(3) closure is structurally derived (§8.6); every established structural result of PST that depends on the embedding choice is verified to hold under both canonical and octonionic embeddings; the order-one condition selects the Yukawa form identically; and the Yukawa-contingency reading of §8.18 is preserved. The only remaining open question, restated as §14.2 the three-generation count, is the structural derivation of $N_{\text{gen}} = 3$ on top of this verified embedding.

8.8. Closure of Conjecture 1: the algebra identity from P1–P3 via Computation 3 and Furey 2014

This subsection closes Conjecture 1 of §14.2: the internal algebra of PST's foundational object is $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, derived from P1–P3 with no further appeal to the Chamseddine-Connes-Marcolli classification. The argument has three pieces, all computationally verified. An earlier route attempt that identified Furey's $\text{Cl}(6)$ as a subalgebra of PST's $\text{Cl}(0, 6)$ (archived) is superseded by the rigorous decomposition delivered in Computations 16 and 19 below.

Piece 1: Computation 3 delivers the chirality grading directly. Computation 3 (`computation_03.py`) showed that on D distinction bits with D odd, the substrate's complementation-and-parity real structure gives $\gamma_K = Z^{\otimes D}$ with $\gamma_K^2 = +I$, sign triple $(\varepsilon, \varepsilon', \varepsilon'') = (+, +, -)$ matching Connes Lorentzian KO-6. Computation 19 §6 (`computation_19.py`) closes the identification of γ_K with the $\text{Cl}(0, 6)$ chirality grading. In the Jordan-Wigner Majorana representation of D bits, the six fermion Majoranas χ_a, χ'_a ($a = 1, 2, 3$) generate $\text{Cl}(0, 6)$ on the $2^D = 8$ -dimensional Hilbert space, and the volume element

$$\omega_{\text{JW}} = \chi_1 \chi'_1 \chi_2 \chi'_2 \chi_3 \chi'_3 \quad (102)$$

satisfies $\omega_{\text{JW}}^2 = -I$. Computation 19 §6 verifies

$$\gamma_K = Z \otimes Z \otimes Z = i \omega_{\text{JW}} \quad (\|\gamma_K - i \omega_{\text{JW}}\| = 0.0 \text{ to machine precision}). \quad (103)$$

The substrate's parity-selected real structure IS the $\text{Cl}(0, 6)$ chirality grading at the operator level, with no convention-dependent step intervening. Hence the substrate naturally provides the chirality projector

$$f = \frac{1}{2}(I + \gamma_K) = \frac{1}{2}(I + i\omega), \quad \bar{f} = \frac{1}{2}(I - i\omega), \quad (104)$$

a rank-4 projector onto the chiral half H_+ of the modal Hilbert space $H_F = \mathbb{C}^8 = H_+ \oplus H_-$.

Piece 2: Spin(6) invariance certifies f as the canonical idempotent. Spin(6) acts on $\text{Cl}(0, 6) \otimes \mathbb{C} \cong M_8(\mathbb{C})$ via the spinor representation, and on $\text{Cl}(0, 6)$ by conjugation through its embedding in $\text{Cl}(0, 6)^{\text{even}}$ as the bivector Lie algebra. Computation 18 (`computation_18.py`) verifies that under this action ω is the unique (up to scalar) Spin(6)-invariant element of positive grade. Concretely, in the Spin(6) decomposition of $\text{Cl}(0, 6)$ by grade

$$\text{Cl}(0, 6) \cong \mathbf{1} \oplus \mathbf{6} \oplus \mathbf{15} \oplus \mathbf{20} \oplus \mathbf{15} \oplus \mathbf{6} \oplus \mathbf{1}, \quad (105)$$

the singlet count is exactly $(1, 0, 0, 0, 0, 0, 1)$ (Computation 18 §6): only the identity (grade 0) and the volume ω (grade 6) are Spin(6)-invariant. Computation 19 §4 then shows that on $H_F = \mathbb{C}^8 = H_+ \oplus H_-$ the Hermitian Spin(6)-invariant endomorphisms form a 2-real-dimensional space spanned by $\{I, \gamma = i\omega\}$. This is Schur's lemma applied to the two irreducible Spin(6) factors $H_+ \cong \mathbf{4}$ and $H_- \cong \bar{\mathbf{4}}$ of $\text{SU}(4) \cong \text{Spin}(6)$. Hence $f = (I + i\omega)/2$ and \bar{f} are the unique (up to chiral swap) Spin(6)-invariant primitive idempotents in $\text{Cl}(0, 6) \otimes \mathbb{C}$. These are exactly Furey's primitive idempotents.¹

Piece 3: Furey 2014 derives A_F as the commutant. Given the primitive idempotent f , Furey 2014 [102] derives:

- a minimal left ideal $\text{Cl}(0, 6) \cdot f$ of real dimension 16, carrying one Standard-Model generation as a representation of $\text{U}(1) \times \text{SU}(2)_L \times \text{SU}(3)$;
- the algebra of operators commuting with this matter representation, the *commutant*,

$$A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C}). \quad (106)$$

The unimodular unitary group of A_F is $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$ (the Standard-Model gauge group) with the matter representation in the correct quantum numbers for one generation.

Conjecture 1, closed. Combining the three pieces:

$$\boxed{\text{P1–P3} + \text{parity selection (Computation 3)} \implies A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})}$$

The chain of derivations is:

1. P1–P3 specify the Boolean substrate, the Bernoulli measure, the asymmetric tension, and the LG modal potential.
2. Computation 3's parity-selected real structure on the odd- D sub-limit delivers $\gamma_K^2 = +I$, sign $(+, +, -)$, KO-6.

¹The construction is from Furey 2014 (*JHEP* 10 (2014) 046, arXiv:1405.4601); cited below Furey [102] in the bibliography as the corresponding three-generation extension.

3. The CAR algebra of the substrate at D bits is $\text{Cl}(0, 2D)$; for $D = 3$ this is the eight-dimensional $\text{Spin}(6)$ spinor representation, on which $M_8(\mathbb{C})$ acts by left multiplication (Computation 16).
4. In the Jordan-Wigner representation, $\gamma_K = i\omega_{\text{JW}}$ exactly (Computation 19 §6, equation (103)). Hence $f = (I + \gamma_K)/2$ is the substrate's chirality projector and coincides with Furey's primitive idempotent.
5. Computation 18 and Computation 19 §4 certify that f is the canonical $\text{Spin}(6)$ -invariant choice (uniqueness up to chiral swap).
6. Furey 2014 [102] derives $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ as the commutant of the matter representation on $\text{Cl}(0, 6) \cdot f$.

Step 4 (equation (103)) is the bridge that closes the gap the referee identified. The derivation is independent of the Chamseddine-Connes-Marcolli classification.

Residual epistemic status. The closure rests on two pieces of standard mathematical content, both well established:

- The CAR algebra of D fermion bits is isomorphic to $\text{Cl}(0, 2D)$ via the Jordan-Wigner Majorana representation.
- Furey 2014's calculation of the commutant of the matter representation on $\text{Cl}(0, 6) \cdot f$ as $\mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$.

Neither is a PST-specific assumption. Under the abstract's *two readings* framing, this closure converts reading (ii) (PST as a derivation of physics from P1–P3 alone) from conjecture-conditional to unconditional for the algebra identity. The remaining open items of §14.2 (the Dirac convergence Sobolev tail, and the three-generation count $N_{\text{gen}} = 3$) are independent of Conjecture 1 and remain open.

8.9. The Standard Model gauge group as substrate automorphisms

The $U(1)_Y$ and $SU(2)_L$ factors are derived from the vacuum topology of Section 8; their combination gives the electroweak gauge group on firm structural footing. Adding the colour factor (Section 8.5) gives the Standard Model gauge group:

$$G_{\text{SM}} = SU(3)_c \times SU(2)_L \times U(1)_Y. \quad (107)$$

The $SU(3)_c$ factor is the unimodular unitary group of the $M_3(\mathbb{C})$ summand of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, so the full G_{SM} is the unimodular unitary group of A_F (Section 14.2). This is the gauge group of the Standard Model as established by Glashow, Weinberg, and Salam [77, 75, 76].

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 [22] (Section 7).

- **Electromagnetism**, the $U(1)_{\text{em}}$ gauge theory from gauging the Goldstone phase mode of $\mathcal{V} \cong S^1 \subset \mathcal{V}_{\text{EW}} \cong S^3$; the photon is the connection on the $U(1)$ bundle over M .
- **The weak force**, the $SU(2)_L$ gauge theory from the $S^3 \cong SU(2)$ vacuum manifold; left-chiral because the threshold crossing is directional; W^\pm and Z masses from the Higgs mechanism on \mathcal{V}_{EW} .
- **The strong force**, the $SU(3)_c$ gauge theory from the $M_3(\mathbb{C})$ colour factor of the internal algebra A_F ; eight massless gluons; confinement from the non-Abelian running coupling.

8.10. Gauge anomaly cancellation as a structural theorem

A gauge theory is quantum-mechanically consistent only if its triangle anomalies vanish [83, 84]: 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.

Important caveat on hypercharges. The $U(1)_Y$ hypercharge assignments used below, $Y = +1/3$ for the quark doublet, $Y = -4/3$ for the right-handed up quark, etc., are taken directly from the Standard Model. PST does not yet derive these values; their derivation from the projection operator Π and the $U(1)_Y$ Goldstone phase is an open problem. The computation below is therefore not a derivation of anomaly cancellation from PST but a demonstration that *given* the SM hypercharges, the PST result $N_c = 3$ is the unique value that makes the theory consistent. The result is: *if* the hypercharges are as in the Standard Model, *then* $N_c = 3$ is forced by internal consistency, and PST fixes $N_c = 3$ independently as the dimension of the colour factor $M_3(\mathbb{C})$ of the internal algebra.

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:

$$[U(1)_Y]^3 = 0 : \quad N_c \left[2\left(\frac{1}{3}\right)^3 + \left(-\frac{4}{3}\right)^3 + \left(\frac{2}{3}\right)^3 \right] + 2(-1)^3 + (2)^3 = -2N_c + 6, \quad (108)$$

$$[SU(2)_L]^2[U(1)_Y] = 0 : \quad N_c \cdot \frac{1}{3} + 1 \cdot (-1) = \frac{N_c}{3} - 1. \quad (109)$$

Both vanish if and only if $N_c = 3$. Since PST fixes $N_c = 3$ as the dimension of the colour factor $M_3(\mathbb{C})$ of the internal algebra A_F (Section 8.15), the following holds as a theorem:

A projected theory with N_c colours is anomaly-free (given the Standard Model)

hypercharge assignments) if and only if $N_c = 3$. PST fixes $N_c = 3$ from the colour factor $M_3(\mathbb{C})$ of the internal algebra, so anomaly cancellation is met at the value the algebra already requires. The Standard Model hypercharge assignments themselves are derived from PST structure in Section 8.11.

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

8.11. $\text{U}(1)_Y$ hypercharges from PST structure

Four structural facts fix all six Standard Model hypercharges without additional input.

(i) *Identification of $\text{U}(1)_Y$.* The PST order parameter $H \in \mathbb{C}^2$ is a complex doublet. Under $\text{SU}(2)_L \times \text{U}(1)_Y$, the $\text{U}(1)_Y$ generator acts as an overall phase rotation $H \rightarrow e^{i\alpha} H$ on both isospin components. The weak hypercharge relation is

$$Q = T_3 + \frac{Y}{2}, \quad (110)$$

where Q is electric charge, $T_3 = \pm \frac{1}{2}$ for an $\text{SU}(2)_L$ doublet, and Y is the $\text{U}(1)_Y$ hypercharge.

(ii) *Higgs hypercharge from photon masslessness.* After the $\text{SU}(2)_L \times \text{U}(1)_Y \rightarrow \text{U}(1)_{\text{em}}$ breaking at $\langle H \rangle = (0, v/\sqrt{2})^T$, the photon is the unique massless gauge boson. This requires the neutral Higgs component H^0 to be electrically neutral: $Q(H^0) = 0$. With $T_3(H^0) = -\frac{1}{2}$ (lower component of the doublet), equation (110) gives

$$Y(H) = 2[Q(H^0) - T_3(H^0)] = 2[0 + \frac{1}{2}] = +1. \quad (111)$$

(iii) *Quark charges from $\text{SU}(3)_c$ confinement.* Colour confinement (Section 8.15) requires all asymptotic states to be colour singlets. For the lightest colour-singlet baryons, the proton ($|uud\rangle$, $Q = +1$) and neutron ($|udd\rangle$, $Q = 0$) give

$$2Q_u + Q_d = 1, \quad Q_u + 2Q_d = 0 \implies Q_u = \frac{2}{3}, \quad Q_d = -\frac{1}{3}. \quad (112)$$

Applied to the left-handed quark doublet $Q_L = (u_L, d_L)$ via equation (110):

$$Y(Q_L) = 2[Q(u_L) - T_3(u_L)] = 2[\frac{2}{3} - \frac{1}{2}] = +\frac{1}{3}. \quad (113)$$

(iv) *Remaining assignments from Yukawa invariance and anomaly cancellation.* The Yukawa couplings $\bar{Q}_L H d_R$, $\bar{Q}_L \tilde{H} u_R$ (where $\tilde{H} = i\sigma_2 H^*$, $Y(\tilde{H}) = -1$), and $\bar{L}_L H e_R$ must each be $\text{U}(1)_Y$

singlets (total charge $Y = 0$):

$$-Y(Q_L) + Y(H) + Y(d_R) = 0 \quad \Rightarrow \quad Y(d_R) = -\frac{2}{3}, \quad (114)$$

$$-Y(Q_L) - Y(H) + Y(u_R) = 0 \quad \Rightarrow \quad Y(u_R) = +\frac{4}{3}, \quad (115)$$

$$-Y(L_L) + Y(H) + Y(e_R) = 0 \quad \Rightarrow \quad Y(e_R) = Y(L_L) - 1. \quad (116)$$

The remaining unknown $Y(L_L)$ is fixed by the $[\text{SU}(2)_L]^2[\text{U}(1)_Y]$ anomaly condition (equation (109)), which requires $N_c Y(Q_L) + Y(L_L) = 0$. With $N_c = 3$:

$$Y(L_L) = -N_c Y(Q_L) = -3 \cdot \frac{1}{3} = -1, \quad Y(e_R) = -2. \quad (117)$$

The complete Standard Model hypercharge assignment is therefore

$$\boxed{Y(Q_L) = +\frac{1}{3}, \quad Y(u_R) = +\frac{4}{3}, \quad Y(d_R) = -\frac{2}{3}, \quad Y(L_L) = -1, \quad Y(e_R) = -2, \quad Y(H) = +1.} \quad (118)$$

Substituting into equation (108) gives $-2N_c + 6 = 0$ at $N_c = 3$, a non-trivial cross-check that no value was inserted by hand.

All six hypercharges are therefore derived consequences of three PST-structural inputs: the order-parameter phase symmetry (identification of $\text{U}(1)_Y$), photon masslessness (fixing $Y(H) = 1$), and $\text{SU}(3)_c$ confinement with $N_c = 3$ (fixing the quark charges and, via anomaly cancellation, the lepton charges). Standard Model hypercharge quantisation is not an independent postulate; it is structurally forced. This closes the gauge group derivation for the bosonic sector. The fermionic sector is established; the colour multiplicity $N_c = 3$ is derived (Section 8.15), and the one-generation fermion-sector representation content is determined explicitly in Section 8.16, matching the hypercharge assignments derived above. The generation count $N_{\text{gen}} = 3$ is *derived* from the V_7 -valued directional structure of P1–P3 and the $\hat{\tau}$ -selection at modal sublimation (§8.17, Computations 46–48): the three quaternionic sub-algebras of \mathbb{O} containing $\hat{\tau}$ are identified with the three generations of matter content. The companion matter-sector observable open in the Standard Model is the Yukawa hierarchy (handled by the no-go theorem below). The Yukawa hierarchy and CKM matrix are addressed by the no-go theorem of §8.18: the substrate’s Bernoulli measure forces the leading-order Yukawa eigenvalues to be S_3 -symmetric and the leading CKM matrix to be the identity, so the observed values are structurally contingent in the precausal configuration $T(C)$ rather than first-principles outputs.

8.12. The Standard Model Lagrangian as leading-order projection

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

$$\mathcal{L}_{\text{SM}} = -\frac{1}{4} F_{\mu\nu}^a F^{a\mu\nu} + \bar{\psi} i \not{D} \psi + |D_\mu H|^2 - \mathcal{F}(H) + \mathcal{L}_{\text{Yukawa}}, \quad (119)$$

where $F_{\mu\nu}^a$ runs over all gauge factors of G_{SM} , $\mathcal{F}(H)$ is the modal potential (63) 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 the coherence scale d_0 . The Yukawa sector $\mathcal{L}_{\text{Yukawa}} = y_{ij}\bar{\psi}_i H\psi_j + \text{h.c.}$ and the fermionic kinetic term require the fermionic sector.

8.13. 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 $\text{SO}(1, 3)$.

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 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 [23] 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. 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 $\mathcal{V}_{\text{EW}} \cong S^3 \cong \text{SU}(2)$ admits half-integer representations; a spin structure on S^3 exists without topological obstruction ($\pi_1(S^3) = 0$). The question of whether the Bernoulli measure gives weight to fermionic sectors is answered in Section 8.14 below: the Boolean $\{0, 1\}$ distinction structure is isomorphic to a fermionic Fock space via the CAR algebra, so fermionic excitations are not an additional postulate but are already present in the substrate.

8.14. Fermionic excitations from the Boolean structure: the CAR algebra identification

The question left open in the preceding subsection is answered by an exact algebraic identification.

The CAR algebra. For each distinction element $a \in D$, define operators c_a, c_a^\dagger on the substrate Hilbert space by

$$c_a^\dagger c_a = n_a, \quad \{c_a, c_b^\dagger\} = \delta_{ab}, \quad \{c_a, c_b\} = 0, \quad (120)$$

where $n_a \in \{0, 1\}$ is the Boolean occupation number of element a . These are the Canonical

Anticommutation Relations (CAR). An explicit realisation is the Jordan-Wigner transformation

$$c_a = \left(\prod_{b < a} (1 - 2n_b) \right) \sigma_a^-, \quad \sigma_a^- = |0\rangle_a \langle 1|_a, \quad (121)$$

which maps the Boolean algebra on $\{0, 1\}^D$ to the CAR algebra exactly, without approximation. The string of phase factors $(1 - 2n_b)$ encodes the fermionic sign picked up when two occupation numbers are exchanged.

Bosonic and fermionic sectors. The full substrate Hilbert space $\mathcal{H}_D = \bigoplus_{k=0}^{|D|} \bigwedge^k \mathbb{C}^{|D|}$ decomposes under the parity operator $(-1)^N$ (where $N = \sum_a n_a$ is the total occupation number) into

$$\mathcal{H}_D^+ = \bigoplus_{k \text{ even}} \bigwedge^k \mathbb{C}^{|D|}, \quad \mathcal{H}_D^- = \bigoplus_{k \text{ odd}} \bigwedge^k \mathbb{C}^{|D|}. \quad (122)$$

States in \mathcal{H}_D^+ (even total occupation) are symmetric under pairwise exchange and project to bosonic fields under Π ; states in \mathcal{H}_D^- (odd total occupation) are antisymmetric and project to fermionic fields.

The Bernoulli measure weights both sectors equally. Under $\mu = \bigotimes_a \text{Bern}(1/2)$, the probability that N is even equals the probability that N is odd: each is exactly $1/2$. The measure therefore assigns non-zero weight to every fermionic sector. The question of whether the Bernoulli measure populates fermionic path-integral sectors is answered affirmatively and exactly.

Spin-statistics connection. The antisymmetry of \mathcal{H}_D^- under exchange is not imposed by hand. It is a structural consequence of the CAR algebra: equations (120) require $c_a c_b = -c_b c_a$ for $a \neq b$, and the Jordan-Wigner string (121) implements this sign automatically. When Π projects \mathcal{H}_D^- to 4D fields, those fields inherit the anticommuting character and satisfy the Dirac equation on (M, g) . The spin-statistics connection (that half-integer spin fields anticommute) is therefore a structural consequence of the Boolean distinction primitive, not a separate postulate.

Status of fermion-sector questions. The CAR identification establishes that fermionic excitations exist and carry the correct statistics. The colour multiplicity $N_c = 3$ is derived via the G_2 decomposition (Section 8.15, equation (124)); the generation count $N_{\text{gen}} = 3$ is *derived* from the V_7 -valued directional structure of P1–P3 and the $\hat{\tau}$ -selection at modal sublimation (§8.17). The *representation content* of the projected fermion sector (which Standard Model multiplets the fermionic Hilbert space \mathcal{H}_D^- maps to under Π at low energy) is determined in Section 8.16 below and matches the §8.3 hypercharge derivation as an independent consistency check. The Yukawa hierarchy and CKM mixing are addressed in §8.18: a no-go theorem (Computation 1) establishes that the substrate’s Bernoulli measure forces leading-order Yukawa eigenvalues to be S_3 -symmetric, so the observed hierarchy is structurally contingent in the precausal configuration $T(C)$ rather than a free parameter to be derived from first principles.

8.15. $SU(3)_c$ from the internal algebra

Colour $SU(3)_c$ is the gauge symmetry of the $M_3(\mathbb{C})$ factor of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$: it is the unimodular unitary group of $M_3(\mathbb{C})$, and the eight gluons arise as the inner fluctuations of the Dirac operator on that factor (Section 14.2).

Octonionic origin. The colour factor descends from the octonions. The exceptional group $G_2 = \text{Aut}(\mathbb{O})$, the automorphism group of the octonions \mathbb{O} , acts irreducibly on the direction algebra $V_7 \cong \mathbb{R}^7$, and the stabiliser of a fixed unit vector in V_7 is $SU(3)$ (Freudenthal–Tits). Under $SU(3) \subset G_2$ the seven directions of V_7 split as

$$\mathbf{7} \longrightarrow \mathbf{1}_{\mathbb{R}} \oplus \mathbf{3}_{\mathbb{C}} \cong \mathbb{R} \oplus \mathbb{C}^3, \quad (123)$$

where $\mathbf{3}_{\mathbb{C}}$ is the colour triplet on which $SU(3)_c$ acts. Equivalently, the complexified octonions $\mathbb{C} \otimes \mathbb{O}$ generate $\text{Cl}(6) \cong M_8(\mathbb{C})$, whose three fermionic ladder operators $\alpha_k = \frac{1}{2}(L_{2k-1} + iL_{2k})$ ($k = 1, 2, 3$) build one generation’s colour triplet (the Gunaydin–Gursey/Furey construction; Section 14.2). The colour multiplicity is the complex dimension of this triplet

$$N_c = \dim_{\mathbb{C}}(\mathbf{3}_{\mathbb{C}}) = 3, \quad (124)$$

the rank of the $M_3(\mathbb{C})$ colour factor.

Colour 3 versus generation 3. The three complex directions of $\mathbf{3}_{\mathbb{C}}$ are the three *colours* of a single generation, not three generations: they are the three colour charges on which $SU(3)_c$ acts. The threefold generation replication is a structurally distinct count: in PST it arises from the V_7 -valued directional structure of P1–P3 and the $\hat{\tau}$ -selection at modal sublimation (§8.17, Computations 46–48). The three quaternionic sub-algebras of \mathbb{O} containing the complex sub-algebra $C_{\hat{\tau}} = \text{span}\{1, \hat{\tau}\}$ are identified with three generations of matter, with $\hat{\tau}$ acting within each generation as the chiral preferred direction. PST therefore derives both $N_c = 3$ (the colour rank inside one generation) and $N_{\text{gen}} = 3$ (the count of generations).

8.16. Fermionic representation content under Π

The Jordan–Wigner identification of Section 8.14 and the $G_2 \rightarrow SU(3) \times U(1)$ decomposition of Section 8.15 together fix the projection Π on the fermionic Hilbert space $\mathcal{H}_{\overline{D}}$ explicitly. Combined with the local-Lorentz spinor bundle $\mathcal{S} = \mathcal{S}^+ \oplus \mathcal{S}^-$ on (M, g) , the *one-generation* fermion content is fully determined.

The one-generation decomposition. The first exterior power $\bigwedge^1 \mathbb{C}^7$ of the substrate fibre, restricted to the chiral sectors of the spinor bundle and refined under $SU(3)_c \times SU(2)_L \times U(1)_Y$, decomposes as

$$\bigwedge^1 \mathbb{C}^7 \Big|_{\Pi}^{1 \text{ gen}} = \begin{aligned} & (\mathbf{3}, \mathbf{2})_{1/3}^L \oplus (\mathbf{3}, \mathbf{1})_{2/3}^R \oplus (\overline{\mathbf{3}}, \mathbf{1})_{-1/3}^R \\ & \oplus (\mathbf{1}, \mathbf{2})_{-1}^L \oplus (\mathbf{1}, \mathbf{1})_{-2}^R \oplus (\mathbf{1}, \mathbf{1})_0^R. \end{aligned} \quad (125)$$

These are, in order, the left-handed quark doublet Q_L , the right-handed up-type singlet u_R , the right-handed down-type singlet d_R , the left-handed lepton doublet L_L , the right-handed

charged lepton e_R , and the right-handed neutrino ν_R .

Hypercharge consistency check. The hypercharges appearing as subscripts in equation (125) reproduce the assignments derived independently in equation (118) via Yukawa-invariance and anomaly cancellation: $Y(Q_L) = +1/3$, $Y(u_R) = +4/3$, $Y(d_R) = -2/3$, $Y(L_L) = -1$, $Y(e_R) = -2$. The right-handed neutrino sits in the singlet $(\mathbf{1}, \mathbf{1})_0$; its appearance is a structural prediction of PST consistent with observed non-zero neutrino masses.

Threefold replication (the generation count is observed, not derived). Nature exhibits three generations. Taking three copies of equation (125) gives the complete Standard Model matter content plus three right-handed neutrinos:

$$\mathcal{H}_D^- \Big|_{\Pi} = \bigoplus_{k=1}^3 \left[Q_L \oplus u_R \oplus d_R \oplus L_L \oplus e_R \oplus \nu_R \right]_k. \quad (126)$$

Total chiral state count: $3 \times (15 + 1) = 48$ Weyl fermions (the standard $45 = 15 \times 3$ SM chiral states plus three right-handed neutrinos). The factor of three here is the *observed* number of generations, not a derived quantity: as discussed in Section 8.15, the three complex directions of $\mathbf{3}_C$ are the three colours of a single generation under the octonion route, so the generation count is not derived from $\mathbf{3}_C$ and remains an open problem. What is derived is the *one-generation* representation content of equation (125); the threefold replication matches observation but its structural origin is not yet established.

What this does and does not predict. Equation (126) predicts the matter content and the hypercharges, both structural. It does *not* predict the Yukawa hierarchy or the CKM matrix. At leading order the three generations in equation (126) are interchanged by the residual SU(3)-flavour symmetry of the G_2 structure (the symmetry that permutes the three complex directions of $\mathbf{3}_C$). Because this symmetry treats the three generations identically, the leading PST Yukawa is degenerate ($y_e = y_\mu = y_\tau = y_0$ at M_*) and the leading PST CKM matrix is the identity. Observation falsifies both leading predictions: the observed lepton mass ratios at M_* are $1 : 206 : 3477$ and the observed CKM mixing is non-trivial in every sector. The hierarchy and the mixing are therefore contingent in the precausal configuration C , not derivable from PST's structural ingredients.

Generation-symmetry no-go. The Bernoulli measure μ on $\{0, 1\}^D$ is the unique probability measure invariant under the full automorphism group $S_D \times \mathbb{Z}_2$ (bit permutations + complementation; paper §3). Generation permutation S_3 embeds as a specific subgroup of S_D : it permutes the three pairs of bits (e_{2k-1}, e_{2k}) assigned to the three generations via the Jordan-Wigner identification. Therefore every observable computed from (D, δ, μ) is automatically S_3 -invariant across generations. This applies in particular to every candidate sub-leading mechanism, Jordan-Wigner phases (vanishing in expectation), multi-instanton sectors (acting on the gauge structure, which is generation-blind by definition), $1/D$ corrections to Mosco convergence (inheriting the full $S_D \times \mathbb{Z}_2$ symmetry), and radion-mediated splittings (acting on the compact fibre direction, shared across generations). Each preserves the SU(3)-flavour symmetry of the

leading structure.

The leading PST Yukawa matrix is therefore necessarily diagonal with equal eigenvalues, and the leading PST CKM matrix is the identity, with no structural mechanism to lift this degeneracy. The observed values must arise from the pre-causal configuration C itself, they are not features of the universal structural form that PST identifies.

SM-sector δm_h^2 cancellation (status sharpened). The structural equal-weighting $N_B = N_F$ of \mathcal{H}_D^+ and \mathcal{H}_D^- under the Bernoulli measure (Section 8.14) is a closed result. However, the numerical Veltman bracket $6\lambda + \frac{3}{2}g_2^2 + \frac{1}{2}g_1^2 - 6y_t^2$ at $M_* = 1573 \text{ GeV}$ evaluates to -3.04 at one loop, an $\sim 8\sigma$ deviation from zero on two-loop precision. PST's structural claim and the numerical Veltman vanishing are therefore not equivalent statements. The structural claim stands; clarifying the precise relation between Bernoulli equal-weighting and the SM-coupling combination is recorded as a residual.

8.17. The three-generation count from the V_7 directional content

The threefold replication of the Standard Model matter content – three generations of fermions with identical quantum numbers – has been the “famously unexplained input” of every noncommutative-geometry Standard-Model framework. In PST, with the V_7 -valued modal order parameter of P3 and the substrate-determined direction $\hat{\tau} = T(D)/|T(D)| \in S^6 \subset V_7$ of P2, the count is a structural theorem.

The structural chain. The direction $\hat{\tau}$ selected at modal sublimation is a unit vector in V_7 . It plays four interlocking roles, all forced by the foundational postulates:

1. *Symmetry breaking.* G_2 , the automorphism group of V_7 (and the symmetry group of the pre-sublimation potential, P3), is broken by $\hat{\tau}$ to its stabiliser, the subgroup of G_2 fixing the unit vector $\hat{\tau}$. This stabiliser is the colour gauge group $SU(3)$: the standard isomorphism $G_2/SU(3) \cong S^6$ means every point on S^6 has $SU(3)$ stabiliser. Verified in Computation 49.
2. *Complex structure on $\hat{\tau}^\perp$.* The 6-dimensional subspace $\hat{\tau}^\perp \subset V_7$ orthogonal to $\hat{\tau}$ inherits a complex structure J given by left-multiplication by $\hat{\tau}$ in the octonion product: $J(v) = \hat{\tau} \times v$ for $v \in \hat{\tau}^\perp$. Computation 49 verifies $J^2 = -I$ on $\hat{\tau}^\perp$, so $\hat{\tau}^\perp \cong \mathbb{C}^3$ carries the $SU(3)$ fundamental representation (the colour content of one fermion generation).
3. *Three quaternionic sub-algebras.* A classical structural fact about \mathbb{O} (verified in Computation 46): exactly three quaternionic sub-algebras of \mathbb{O} contain the complex sub-algebra $C_{\hat{\tau}} = \text{span}\{1, \hat{\tau}\}$. For $\hat{\tau} = e_1$, these are

$$Q_1 = \{1, e_1, e_2, e_3\}, \quad Q_2 = \{1, e_1, e_4, e_5\}, \quad Q_3 = \{1, e_1, e_6, e_7\},$$

one per Fano-plane line through $\hat{\tau}$. The three sub-algebras partition $\hat{\tau}^\perp$ (the 6-dim perpendicular space) into three disjoint 2-dim pairs. The count of 3 is G_2 -symmetric: it does not depend on which unit vector in S^6 is selected.

4. *Three Cl(0,2) chiral splittings.* Each Q_k generates a Cl(0,2) sub-Clifford structure on \mathbb{R}^8 via left multiplication by its two non- $\hat{\tau}$ generators. Computation 47 verifies each Q_k delivers a 4 + 4 chiral split with volume $\omega_k^2 = -I$, structurally identical across the three Q_k . This matches exactly the observed pattern of three generations with identical Standard-Model quantum numbers.

Three disjoint generation sectors. The chain-algebra space $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O} \cong \mathbb{R}^{64}$, which carries the matter representation under PST's Cl(0,6) chirality + Furey 2014 chain (§8.8), decomposes under the $\hat{\tau}$ -selection as

$$\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O} = \underbrace{\mathbb{C} \otimes \mathbb{H} \otimes \text{span}\{1, \hat{\tau}\}}_{\text{core, dim=16}} \oplus \bigoplus_{k=1}^3 \underbrace{\mathbb{C} \otimes \mathbb{H} \otimes \text{span}\{e_{a_k}, e_{b_k}\}}_{G_k, \text{dim=16}}. \quad (127)$$

The decomposition is verified in Computation 48. The arithmetic is exact: $16 + 3 \times 16 = 64$. The pairwise intersections $G_i \cap G_j = 0$ for $i \neq j$; the three generation-specific sectors are mutually disjoint as real subspaces of \mathbb{R}^{64} . The 16-dimensional **core** carries the gauge structure (the algebra A_F acting on its own minimal ideal); it is generation-blind. Each G_k carries one generation's matter content.

The per-generation Hilbert space is $\mathcal{H}_k = \text{core} \oplus G_k$ of real dimension 32. Its chiral half (under the extended chiral projector $f = (I + i_{\mathbb{C}} \otimes I_{\mathbb{H}} \otimes \omega_6)/2$ of Computation 45) has dimension 16, matching exactly the Standard-Model count of 15 chiral fermions plus one right-handed neutrino per generation (§8.16). The chiral structure is identical across the three \mathcal{H}_k , recovering the observed pattern of three generations with identical quantum numbers.

Colour and generation as structurally distinct. The number 3 appears in PST in two structurally distinct guises:

- *Colour* $N_c = 3$: the rank of the $M_3(\mathbb{C})$ factor of A_F (§8.5). It arises *within* a single generation, from the $SU(3) \subset G_2$ stabiliser of $\hat{\tau}$.
- *Generations* $N_{\text{gen}} = 3$: the count of quaternionic sub-algebras of \mathbb{O} containing the complex sub-algebra $C_{\hat{\tau}}$. It arises *across* generations, from the $\hat{\tau}$ -direction selection in V_7 .

Both counts emerge from the same $\hat{\tau}$ -selection event, but they parameterise different structural data: colour is internal to each generation's \mathcal{H}_k via the $M_3(\mathbb{C})$ factor of A_F ; generation 3 counts the \mathcal{H}_k themselves.

Status. The structural derivation of $N_{\text{gen}} = 3$ via equation (127) is verified numerically in Computations 46–48. The remaining concrete tasks are: (i) demonstrate that each G_k carries the explicit $SU(3) \times SU(2) \times U(1)$ representation content matching the one-generation decomposition of equation (125); (ii) establish consistency of the $\hat{\tau}$ -determination $\hat{\tau} = T(D)/|T(D)|$ with the Mosco-limit projection $\Phi : S \rightarrow G$. Both are concrete operator-algebra tasks of the same character as the section-round-trip and Dirac-aware-lift closures of §7.7.

8.18. Yukawa hierarchy and CKM mixing as contingent in $T(C)$: a structural no-go

The observed Yukawa hierarchy (charged-lepton masses in the ratio $1 : 206 : 3477$ at M_*) and the non-trivial CKM mixing matrix are *not* structural predictions of PST. They are contingent on the precausal configuration C — specifically, on the asymmetric-tension distribution $T(C)$ that produced this universe. This is a precise statement of PST’s scope, established as a theorem in Computation 1 (`computation_01.py`).

No-go theorem (substrate level). The Bernoulli product measure $\mu = \bigotimes_a \text{Bern}(1/2)$ on $\mathcal{P}(D)$ is invariant under the full permutation group $S_D \times \mathbb{Z}_2$ acting on the D substrate bits. Generation permutation $S_3 \subset S_D$ (permuting the bit-triples assigned to the three generations) is therefore a symmetry of any observable computed from (D, δ, μ) alone. Computation 1 verifies this numerically for the leading proposed generation-asymmetric mechanisms: the Yukawa overlap integral, the multi-instanton Atiyah–Singer index per representation, and the Jordan-Wigner expectation in the even-parity sector are all S_3 -invariant. Therefore the leading PST Yukawa coupling is diagonal with equal eigenvalues, the leading CKM matrix is the identity, and the observed hierarchy and mixing must enter from outside the structural ingredients. The contingency is located precisely in $T(C)$.

Spectral-triple extension. The no-go survives the Connes-Chamseddine spectral-triple framework via inner-fluctuation analysis. Since A_F acts on H_F as $a^{(1)} \otimes I_{\text{gen}}$ with no generation labels, an inner fluctuation $\omega = \sum a_i [D, b_i]$ inherits the generation structure of D_F unchanged; it cannot lift a degeneracy absent from D_F nor remove one present. The generation pattern is locked in D_F , which is a contingent input. This extension is a direct corollary of the bimodule structure (Computation 1 §7).

Quantitative contingency. The observed flavour parameters (approximately twenty for Dirac neutrinos: nine Yukawa eigenvalues, four CKM parameters, seven PMNS) consume about 31% of the $2^D = 64$ real degrees of freedom available in $T(C)$ at $D = 6$ substrate bits (Computation 1 §8). PST is therefore neither over-determined nor under-equipped to encode the observed flavour pattern: there is ample structural room for the contingency, with a definite count of how much is consumed.

This is the honest scope-statement: PST does not derive contingent parameter values; it derives the structural form within which those parameters live. In good company, this matches the noncommutative-geometric Standard Model of Connes-Chamseddine, in which the Yukawa eigenvalues are inputs to D_F rather than outputs.

8.19. Chirality of $\text{SU}(2)_L$ as a structural consequence

The chirality of the weak gauge group (that $\text{SU}(2)_L$ acts on left-handed doublets only) is a structural consequence of PST’s substrate combined with the directed modal threshold of paper §3. This is the chiral-bimodule support property of the Connes-Chamseddine finite triple, established for PST in Computations 7, 8, and 11.

The substrate's three-sphere factor S^3 in $M = \mathbb{R} \times S^3$ carries an isometry algebra $\mathfrak{so}(4) = \mathfrak{su}(2)_L \oplus \mathfrak{su}(2)_R$, equivalently the algebra of quaternion left and right multiplications on \mathbb{H} . This furnishes *two* candidate $\mathfrak{su}(2)$ representations on the internal Hilbert space H_F , one acting on each chirality block of γ_F . Both have clean block support and the two are exact parity conjugates of each other ($P \cdot H_A \cdot P^{-1} = H_B$ verified numerically; Computation 11, `computation_11.py`). Together with Computation 7's verification that the order-one condition selects the Yukawa form given the chiral bimodule, the chirality reduces to a single physical question: which of the two parity-conjugate $\mathfrak{su}(2)$ embeddings is realised in nature?

The answer is the parity-violation mechanism PST already commits to in §3: the directed modal threshold $\varepsilon > 0$ orients the S^3 volume form, fixing one factor as the physical $\mathrm{SU}(2)_L$. The same mechanism that produces parity violation selects the chiral bimodule. Chirality of $\mathrm{SU}(2)_L$ is therefore structural, conditional only on the threshold-direction parity argument PST commits to elsewhere.

8.20. 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 gauge structure from the substrate's internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$. The algebra itself is derived from P1–P3 in §8.8 (Computations 3, 21, and Furey 2014), not posited; the gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ is then the unimodular unitary group of A_F , with gauge bosons and the Higgs arising as inner fluctuations of the Dirac operator. The colour group $\mathrm{SU}(3)_c$ is the $M_3(\mathbb{C})$ factor of A_F , and the octonionic origin of the internal algebra (via the direction algebra $V_7 \mathrm{Im}(\mathbb{O})$ on which G_2 acts) supplies one generation's colour and electric-charge content. A structural consistency requirement (gauge anomaly cancellation, which demands that the number of quark colour charges be exactly three) is not an accident but a theorem conditional on $N_c = 3$: the same $N_c = 3$ that the algebra delivers also makes the quantum theory self-consistent. The generation count $N_{\mathrm{gen}} = 3$ is *not* derived: under the octonion route the three complex directions are spent on colour, so the threefold generation replication is a separate open question (the three-generation count of §14.2). The relation $Z^2 \approx e^{-1}$ between the PST field-normalisation factor and the Standard Model coupling enters as a one-loop numerical near-coincidence, noted but not derived from first principles.

9. Physical Predictions

9.1. Structural and specific 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 this 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.

9.2. The Born rule from the Bernoulli measure

A foundational question that Section 8 noted explicitly as open is the derivation of the Born rule: the identification of $|\psi|^2$ with measurement probability. This subsection provides a structural derivation rooted in the Bernoulli measure and the Hilbert-space structure established in Section 3.

The PST Hilbert space. The gradient structure derivation of Section 3 established that the space $L^2(\mathcal{P}(D), \mu)$ of square-integrable complex functionals on configuration space, with the Bernoulli measure $\mu = \bigotimes_a \text{Bern}(1/2)$, is a separable Hilbert space. The inner product is

$$\langle \phi, \psi \rangle = \int_{\mathcal{P}(D)} \phi^*(C) \psi(C) d\mu(C), \quad (128)$$

and the order parameter $\psi \in L^2(\mathcal{P}(D), \mu)$ is a unit vector (after normalisation). This is the quantum state space of PST, derived from the Bernoulli measure rather than postulated.

Physical observables as projectors. A physical observable \hat{O} with outcomes $\{o_i\}$ and corresponding eigenspaces $\{\mathcal{H}_i\}$ defines an orthogonal decomposition $L^2(\mathcal{P}(D), \mu) = \bigoplus_i \mathcal{H}_i$. The corresponding projection operators $P_i : L^2 \rightarrow \mathcal{H}_i$ satisfy $P_i P_j = \delta_{ij} P_i$ and $\sum_i P_i = \mathbf{1}$.

The $|\psi|^2$ probability from the Bernoulli measure. The Bernoulli measure μ assigns probability $\mu(S)$ to each measurable subset $S \subseteq \mathcal{P}(D)$. Near the threshold, the physical probability of obtaining outcome o_i is the rate at which configurations in the eigenspace \mathcal{H}_i cross the modal threshold τ . The threshold-crossing rate for a configuration C is proportional to the modal tension excess $|a_{LG}||\psi(C)|^2$ (from the linear stability analysis of Section 3): configurations where ψ has large amplitude cross the threshold more frequently, and the crossing rate is proportional to $|\psi|^2$, not to $|\psi|$ or $|\psi|^3$. The probability weight associated with configuration

C crossing at outcome o_i is therefore $|\psi(C)|^2 \cdot \mathbf{1}_{\mathcal{H}_i}(C)$, giving:

$$P(\text{outcome } o_i \mid \psi) = \frac{\int_{C:\psi(C)\in\mathcal{H}_i} |\psi(C)|^2 d\mu(C)}{\int_{\mathcal{P}(D)} |\psi(C)|^2 d\mu(C)} = \frac{\|P_i\psi\|^2}{\|\psi\|^2} = \|P_i\psi\|^2 \quad (129)$$

for unit-normalised ψ . This is the **Born rule**.

Gleason's theorem closes the argument. The probability assignment $P \mapsto \|P\psi\|^2$ satisfies all probability axioms and, crucially, it is *additive on orthogonal projectors*: $P(P_i \oplus P_j) = \|P_i\psi\|^2 + \|P_j\psi\|^2$ for orthogonal P_i, P_j . This additivity is guaranteed by the linearity of the Bernoulli integral: $\int_{S_i \cup S_j} |\psi|^2 d\mu = \int_{S_i} |\psi|^2 d\mu + \int_{S_j} |\psi|^2 d\mu$ for disjoint S_i, S_j . By Gleason's theorem [52] applied to the separable Hilbert space $L^2(\mathcal{P}(D), \mu)$ (dimension ≥ 3 for $|D| \geq 2$), the *unique* non-contextual probability measure on the orthogonal projectors of a Hilbert space is $P \mapsto \|P\psi\|^2$. The Born rule is therefore not an additional postulate of PST but a consequence of the Bernoulli measure and the Hilbert-space structure of $L^2(\mathcal{P}(D), \mu)$.

What this establishes and what remains open. The derivation above establishes the Born rule for measurements that decompose into orthogonal eigenspaces in $L^2(\mathcal{P}(D), \mu)$. What it does not yet supply is the full measurement postulate (collapse of the wave function) and the treatment of the fermionic sector. The canonical commutation relations are recovered in the following subsection, given the identification of the projected Boolean gradient with ∂_μ .

9.3. Canonical commutation relations from the Boolean gradient

The projected Hilbert space $\mathcal{H} = L^2(M, d^4x)$ carries a natural position operator:

$$(\hat{x}_\mu \psi)(x) = x_\mu \psi(x), \quad (130)$$

defined as multiplication by the spacetime coordinate x_μ . The momentum operator is the image of the Boolean gradient under the projection Π :

$$\hat{p}_\mu = -i\hbar \Pi(\partial_\mu^{\text{B}}), \quad (131)$$

where ∂_μ^{B} is the Boolean (Efron–Stein) gradient (equation (17)) in the μ -direction.

Continuum limit of the Boolean gradient. For a smooth function $f \in L^2(M)$ pulled back to $L^2(\mathcal{P}(D), \mu)$ via Π , the Boolean central difference acting on a single substrate element a at position x_a is:

$$(\partial_a^{\text{B}}\psi)(C) = \psi(C \cup \{a\}) - \psi(C \setminus \{a\}) \xrightarrow{|dx_a| \rightarrow 0} 2 \partial_{x_\mu} f(x) \cdot \hat{e}_\mu, \quad (132)$$

where \hat{e}_μ is the unit vector in the μ -direction and the factor 2 arises because adding and removing a displace the projected centre $x(C)$ by $+dx_a$ and $-dx_a$ respectively. Including the factor $\frac{1}{2}$ implicit in the projection gives $\Pi(\partial_\mu^{\text{B}}) = \partial_\mu$, so that $\hat{p}_\mu = -i\hbar \partial_\mu$ acts on $L^2(M, d^4x)$

as the standard momentum operator.

Derivation. Acting on any $\psi \in \mathcal{H}$:

$$\begin{aligned} [\hat{x}_\mu, \hat{p}_\nu] \psi &= -i\hbar(x_\mu \partial_\nu \psi - \partial_\nu(x_\mu \psi)) \\ &= -i\hbar(x_\mu \partial_\nu \psi - \delta_{\mu\nu} \psi - x_\mu \partial_\nu \psi) \\ &= i\hbar \delta_{\mu\nu} \psi. \end{aligned} \tag{133}$$

The canonical commutation relations $[\hat{x}_\mu, \hat{p}_\nu] = i\hbar \delta_{\mu\nu}$ follow once $\hat{p}_\mu = -i\hbar \partial_\mu$ is identified on $L^2(M, d^4x)$. The product rule step (second line) is then trivial. The substantive step is the continuum limit (132): the identification of the Boolean gradient with the standard momentum operator. The heuristic argument above is made rigorous in the following paragraph via strong resolvent convergence.

Rigorous continuum limit via strong resolvent convergence. The Mosco-convergence result of Section 7.6 (equation (46)) makes the limit (132) rigorous. Strong resolvent convergence of the rescaled Boolean Laplacian to Δ_g implies that for any $f \in H^1(M)$ and $\lambda > 0$, the Boolean gradient $|D|^{-1/2} \nabla^B(\Pi^\dagger f)$ converges in $L^2(\mathcal{P}(D), \mu)$ to the Riemannian gradient $\nabla_g f$. Under Π , this maps $\partial_a^B \mapsto \partial_\mu$ in $L^2(M, d^4x)$, so $\hat{p}_\mu = -i\hbar \partial_\mu$ is the strong limit of the projected Boolean gradient. The operator $-i\hbar \partial_\mu$ is essentially self-adjoint on $C_c^\infty(M) \subset L^2(M, d^4x)$ by standard elliptic theory [93]; its closure is the unique self-adjoint extension on the Sobolev domain $H^1(M)$. The continuum-limit identification $\Pi(\partial_\mu^B) = \partial_\mu$ is therefore not merely heuristic: the Boolean gradient converges to the standard momentum operator in the strong resolvent topology, and the resulting \hat{p}_μ is self-adjoint and densely defined. The canonical commutation relations are thereby established.

9.4. Quantum Fluctuations

Near the modal threshold, where $\varepsilon(C) = T(C) - \tau \rightarrow 0$, 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:

$$\delta g_{\mu\nu} \sim \eta_{\mu\nu} \cdot \frac{\delta T}{2|T - \tau_{\text{null}}|^{1/2}} \Big|_{T \rightarrow \tau_{\text{null}}} \tag{134}$$

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 (29) 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.

9.5. 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 (29) and the definition of the Ricci scalar:

$$R(\Phi(C)) \propto \Pi(\nabla_C^2 T(C)) \quad (135)$$

Curvature is the projection of the configuration-space Laplacian of asymmetric tension onto the instantiated manifold. The Laplacian ∇_C^2 acts on the discrete power set $\mathcal{P}(D)$; its continuum realisation via Π is an open problem. 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 [45] 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 $\nabla_C^2 T(C) \rightarrow \infty$ 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 [46, 47] acquires a precausal interpretation as the thermal fluctuation of near-threshold configurations at the horizon, consistent with the thermodynamic derivation of Jacobson [19] but now grounded in the precausal structure of Π rather than in a phenomenological entropy argument.

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

Equation (39) 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:

$$\nabla^2 \bar{h}_{00} = -16\pi G \rho \tag{136}$$

Setting $\bar{h}_{00} = -4\phi$, where ϕ is the Newtonian gravitational potential ($g_{00} \approx -(1 + 2\phi)$), equation (136) reduces to:

$$\nabla^2 \phi = 4\pi G \rho \tag{137}$$

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 (39) 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 .

9.7. 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.

9.8. 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]:

$$P_{\text{standard}}(d) = -\frac{\pi^2 \hbar c}{240 d^4} \quad (138)$$

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 $\rho : D \rightarrow M$ 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:

$$P_{\text{PST}}(d) = P_{\text{standard}}(d) \left[1 + \xi \left(\frac{d_0}{d} \right)^2 + \mathcal{O} \left(\left(\frac{d_0}{d} \right)^4 \right) \right] \quad (139)$$

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

$$\delta P_{\text{PST}}(d) = - \frac{\xi \pi^2 \hbar c d_0^2}{240 d^6} \quad (140)$$

The d^{-6} scaling is a necessary consequence of any single-length-scale, isotropic short-distance modification of the vacuum: the leading even power of d_0/d multiplied by the standard d^{-4} Casimir term gives d^{-6} regardless of the specific theory. Absence of a d^{-6} correction would falsify PST; presence of a d^{-6} correction is consistent with PST but does not select it from other minimal-length or EFT-cutoff models. The quantity specific to PST is the *amplitude* ξd_0^2 : $\xi = 90/\pi^2$ is fixed exactly by the kernel K of the projection operator (below), while the scale d_0 is the Landau-Ginzburg coherence length, set by the LG dynamics.

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

A careful reader will observe that equation (140) contains two quantities whose precise values determine the amplitude: the coefficient ξ and the scale d_0 . The situation is as follows.

The discreteness scale d_0 is the Landau-Ginzburg coherence length; its value is set by the LG vacuum structure and is not pinned by the three foundational postulates alone. The coefficient ξ , by contrast, is determined exactly by the Gaussian moments of the PST projection kernel K :

$$\xi = \frac{15}{\pi^2} \frac{\int_0^\infty e^{-x^2/(4d_0^2)} x^4 dx}{d_0^2 \int_0^\infty e^{-x^2/(4d_0^2)} x^2 dx} = \frac{15}{\pi^2} \times 6 = \frac{90}{\pi^2} \approx 9.12, \quad (141)$$

Dimensional verification: ξ is a pure number. Substituting $u = x/(2d_0)$ evaluates the integrals exactly:

$$\int_0^\infty e^{-x^2/(4d_0^2)} x^4 dx = 12\sqrt{\pi} d_0^5, \quad \int_0^\infty e^{-x^2/(4d_0^2)} x^2 dx = 2\sqrt{\pi} d_0^3. \quad (142)$$

Hence the numerator carries dimension d_0^5 (length⁵) and the denominator carries $d_0^2 \times d_0^3 = d_0^5$ (length⁵), so the ratio is dimensionless: $12\sqrt{\pi} d_0^5 / (d_0^2 \times 2\sqrt{\pi} d_0^3) = 6$, and $\xi = 15 \times 6 / \pi^2 = 90/\pi^2$. The d_0^2 in the denominator of equation (141) is not “dangling”: it is the factor that cancels the d_0^5 from the x^4 moment against the d_0^3 from the x^2 moment, leaving a pure number. This

cancellation was verified symbolically (Computations §verify_track_b).

The coefficient $\xi = 90/\pi^2 \approx 9.12$ is therefore fixed exactly; the amplitude ξd_0^2 is parameter-free in its coefficient, with the scale d_0 the Landau-Ginzburg coherence length.

What does *not* depend on ξ or d_0 is the power-law *exponent*. The exponent -6 in $\delta P_{\text{PST}} \propto d^{-6}$ 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 Lifshitz-theory prediction (which already accounts for finite conductivity, surface roughness, and thermal fluctuations to high precision), show a residual that scales as d^{-6} ? If the residual is consistent with zero within experimental uncertainty, PST is constrained. If it shows a different power law that cannot be attributed to known corrections, PST is falsified. A d^{-6} residual is a *necessary* condition for PST at this order, but not sufficient: other single-length-scale modifications of the vacuum also predict d^{-6} [15]. At the first layer PST therefore makes a categorical prediction about *the form of the correction* (the exponent), not about its amplitude.

The *second layer*, conditioned on observing d^{-6} scaling, is quantitative. $\xi = 90/\pi^2$ is fixed exactly by the projection-kernel moments; the amplitude ξd_0^2 contains a single free parameter d_0 , the Landau-Ginzburg coherence length, which is *not derived from P1–P3* and which existing 160–750 nm precision Casimir data constrains only by an upper bound, $d_0 < 16$ nm (§14.3). A measured residual at d^{-6} with amplitude A therefore determines $d_0 = \sqrt{A/\xi}$ from the data rather than confirming a parameter-free prediction. Inconsistency between the inferred d_0 from different experimental separations *would* falsify PST at the second layer; a single amplitude measurement at the precision required to extract a d^{-6} residual constitutes a determination of d_0 , not a sharp test.

Honest comparison with parameter-rigid tests: the gravitational-wave tests of General Relativity constrain a specific waveform morphology in which both the exponent (post-Newtonian structure) and the amplitude (Bayesian fit to source parameters) are predicted up to a finite list of source-parameter inputs (masses, spins). PST’s Casimir test is exponent-rigid but not amplitude-rigid until d_0 is independently fixed; in this sense PST’s current falsifiability is intermediate between a parameter-free theory and a fit, closer to the parameter-free end at the exponent layer and closer to the fit end at the amplitude layer. The asymmetry is real and is acknowledged in the open-research register (§14.2): deriving d_0 from the postulates, or

anchoring it to another PST observable, is required to upgrade the second layer to a sharp parameter-free test.

The d^{-6} scaling is separable from the standard d^{-4} Casimir dependence by fitting measurements at multiple separations. The correction budget at relevant separations is discussed in detail in Section 10.10. In brief: finite-conductivity corrections enter via the Lifshitz formula as an effective d^{-5} correction at leading order in $\lambda_{\text{plasma}}/d$, and this Lifshitz-corrected Casimir force is well-known and can be subtracted with high precision. The PST d^{-6} signal would appear as a residual after this subtraction. Surface roughness introduces corrections at the spatial-frequency scale of the roughness spectrum, with a known power-law dependence; thermal corrections have a distinct temperature dependence isolable by varying plate temperature. The subtraction strategy is the standard one used in current precision Casimir experiments [15]; the challenge is the accuracy required to resolve a $\sim 2\%$ residual at $d = 50$ nm after subtracting Lifshitz corrections that are themselves of order 100% at this separation.

9.10. 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:

$$d_0 < \frac{16 \text{ nm}}{\sqrt{|\xi|}} \quad (143)$$

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:

$$\frac{P(d_1)/P_{\text{standard}}(d_1)}{P(d_2)/P_{\text{standard}}(d_2)} - 1 \propto \left(\frac{d_0}{d_1}\right)^2 - \left(\frac{d_0}{d_2}\right)^2 \quad (144)$$

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

What the theory fixes without free parameters is the *form* of the correction: the d^{-6} scaling and the coefficient $\xi = 90/\pi^2 \approx 9.12$ (equation (141)). The overall amplitude scales as $(d_0/d)^2$, with d_0 the Landau-Ginzburg coherence length of the modal condensate (Section 10); its absolute magnitude is a condensate property, naturally in the nanometre range and many orders of magnitude above ℓ_P , rather than a number the three postulates pin down. For a coherence length of order a few nanometres the correction is at the edge of current precision near $d = 160$ nm, reaches several percent near $d = 50$ nm, within reach of next-generation experiments, and becomes order-unity below $d = 20$ nm. The Casimir measurement is accordingly a genuine predicted detectable signal, not merely a formal bound, with its precise size set by d_0 .

9.11. 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 and the Casimir Prediction

Conditionality. The parameter-free content of this section (the d^{-6} Casimir scaling and its coefficient $\xi = 90/\pi^2$) depends only on the moments of the projection kernel and is independent of the magnitude of d_0 . The absolute size of the Casimir signal at a given separation scales as $(d_0/d)^2$ and is therefore set by the coherence length d_0 , whose value is a property of the modal condensate rather than an output of (D, δ, μ) alone. The mathematical content is sound within that scope.

10.1. The default assumption that must be dropped

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: d_0 is not the microscopic grain of the substrate but the coherence length of the modal condensate, which near a second-order transition sits many orders of magnitude above the grain. This places d_0 in the nanometre range, far above ℓ_P , so that the Casimir correction is a genuine predicted detectable signal, not a formal statement about an unobservably small parameter.

10.2. 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:

$$\xi_{\text{BCS}} = \frac{\hbar v_F}{\pi \Delta} \quad (145)$$

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 $\xi_{\text{BCS}} \sim 10^3 a$ 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 $\rho : D \rightarrow M$ 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:

$$F[\psi] = \int d^4x \left[\frac{1}{2} c_{LG} |\nabla\psi|^2 + a_{LG} |\psi|^2 + b_{LG} |\psi|^4 \right] \quad (146)$$

the coherence length of ψ is [8]:

$$d_0 = \xi_\psi = \sqrt{\frac{c_{LG}}{2|a_{LG}|}} \quad (147)$$

Near the modal threshold, $a_{LG} \propto (T/\tau - 1)$, so $\xi_\psi \rightarrow \infty$ as $T \rightarrow \tau^+$: 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, a property of the modal condensate rather than a number the postulates fix.

10.3. Modal sublimation and the substrate scales

The functor $\Phi : \mathcal{S} \rightarrow \mathcal{G}$ 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. The remaining substrate structure is carried by the finite noncommutative internal factor (Section 14.2), which has no spatial extent: it is the algebraic home of the gauge and matter structure, not a set of compact spatial dimensions.

Two length scales characterise the instantiated theory. The first is the Landau-Ginzburg

coherence length of the order parameter,

$$d_0 = \sqrt{\frac{c_{LG}}{2|a_{LG}|}}, \quad (148)$$

which sets the substrate discreteness scale that enters the Casimir correction. The second is the modal scale $\ell_* = \hbar/M_*$, the length associated with the energy M_* at which the modal condensate forms (identified below with $M_* = 4\pi m_h$). With the Higgs identification $r_0 = v$ and $|a_{LG}| = m_h^2/4$, the coherence length d_0 is fixed by the Landau-Ginzburg vacuum structure rather than by any compactification, and its numerical value is not determined by the three postulates alone.

10.4. 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.

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 (64) 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 [85, 86], with no additional input. The modal scale $M_* = 4\pi m_h \approx 1573$ GeV is a fourth and separate quantity (Section 7.20): it is not the vev and not the radial-mode mass but the threshold energy $\hbar c/\ell_*$ at which the full modal structure becomes manifest. The Higgs mass and the vev are the low-energy observables of the modal condensate. All three are mutually consistent through $b = \lambda/4$ and $v < M_*$.

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_* is settled. $M_* = 4\pi m_h \approx 1573$ GeV is derived as a structural prediction of PST (Section 7.20) from the requirement that the one-loop LG quartic matches the SM Higgs self-coupling at threshold. This derivation uses $m_h = 125.20$ GeV as the sole measured input and requires no free parameters beyond the measured input m_h . The structural prediction throughout the rest of the paper is $M_* = 4\pi m_h \approx 1573$ GeV.

The substrate grain length is then $\ell_* = \hbar c/M_*$. The Casimir correction's power-law exponent and its coefficient $\xi = 90/\pi^2$ are structural predictions; the onset scale d_0 is the Landau-Ginzburg coherence length, set by the LG dynamics. Taking $M_* = 4\pi m_h \approx 1573$ GeV:

$$\ell_* = \frac{\hbar c}{M_*} = \frac{197.3 \text{ MeV} \cdot \text{fm}}{1573 \text{ GeV}} \approx 1.253 \times 10^{-19} \text{ m} \quad (149)$$

10.5. The substrate coherence scale

The Casimir correction below is governed by the substrate coherence length d_0 of equation (148). Here d_0 is set by the Landau-Ginzburg vacuum structure and is not fixed in value by the three foundational postulates alone; accordingly the Casimir correction is a parameter-free prediction in its power-law *exponent*, with the onset scale d_0 left to be determined by the Landau-Ginzburg dynamics or by experiment.

10.6. Casimir detectability at each separation regime

Evaluating at a representative coherence length $d_0 \approx 4.92$ nm (the theory fixes the d^{-6} form and the coefficient $\xi = 90/\pi^2 \approx 9.12$ exactly, but not the magnitude of d_0), the Casimir correction of equation (140) gives a concrete signal at each accessible plate separation.

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

$$\left(\frac{d_0}{d}\right)^2 \Big|_{d=160 \text{ nm}} = \left(\frac{4.92}{160}\right)^2 \approx 9.45 \times 10^{-4} \quad (150)$$

The PST correction with $\xi = 90/\pi^2$ is $9.12 \times 0.094\% \approx 0.86\%$. This lies at the edge of the current 1% precision floor; the theory is consistent with all existing data, and the correction may become detectable with current-generation instruments.

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

$$\left(\frac{d_0}{d}\right)^2 \Big|_{d=50 \text{ nm}} = \left(\frac{4.92}{50}\right)^2 \approx 9.68 \times 10^{-3} \quad (151)$$

The correction with $\xi = 90/\pi^2$ is $9.12 \times 0.97\% \approx 8.8\%$, well above the 1% current floor and unambiguously detectable at next-generation sensitivity.

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

$$\left(\frac{d_0}{d}\right)^2 \Big|_{d=20 \text{ nm}} = \left(\frac{4.92}{20}\right)^2 \approx 0.0606 \quad (152)$$

A correction of $9.12 \times 6.1\% \approx 55\%$ 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 the same representative $d_0 \approx 4.92$ nm explicitly into equation (140):

$$\delta P_{\text{PST}}(d) = -\frac{\xi \pi^2 \hbar c \times (4.92 \text{ nm})^2}{240 d^6} \quad (153)$$

The coefficient $\xi = 90/\pi^2 \approx 9.12$ is derived exactly from the Gaussian moments of the substrate weight function; see equation (141). Measuring the d^{-6} component at two separations via the ratio test (144) 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 (144)) 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 (143). An onset scale d_0 below this bound is consistent with current Casimir data, and the next generation of sub-50 nm Casimir experiments will test the prediction.

The d^{-6} correction structure is independent of the numerical value of d_0 : its power-law exponent and its coefficient $\xi = 90/\pi^2$ are structural, while the onset scale d_0 is the Landau-Ginzburg coherence length set by the LG dynamics. A measured d^{-6} residual at sub-50 nm separations would support the substrate-discreteness prediction; its absence at any separation would falsify it. The modal scale $M_* = 4\pi m_h \approx 1.573$ TeV is a separate, collider-testable prediction: PST expects new physics associated with the modal condensate near this scale.

10.8. Experimental status of M_* and the ADD comparison

PST is superficially analogous to the ADD model [16] only in involving structure beyond the four macroscopic dimensions. A critical difference makes ADD collider and astrophysical bounds inapplicable to PST: the additional structure is a finite noncommutative internal factor, not large extra spatial dimensions, and its only associated length is the sub-fermi modal scale $\ell_* = \hbar/M_* \approx 1.25 \times 10^{-19}$ m, far from the millimetre-to-micron range ADD requires for its gravity-dilution mechanism. The ADD bounds on the fundamental scale (LHC mono-jet/mono-photon ($M_* \gtrsim 5\text{--}8$ TeV), SN1987A and neutron-star KK-graviton emission) are all derived for macroscopic extra dimensions, with KK gravitons coupling to the full stress-energy tensor at tree level with the four-dimensional Newton constant. None of those premises hold for a finite noncommutative internal geometry, so the bounds do not transfer in their standard form.

What remains is a genuine empirical-consistency question, not a derived constraint: whether $M_* = 4\pi m_h \approx 1.573$ TeV is consistent with the existing data once the PST-specific coupling structure (set by the projection operator Π) is computed. That computation has not been performed, and the question is listed among the open items (Section 14.2, item (E)). The coherence length d_0 that sets the absolute size of the Casimir signal is a property of the modal condensate; its magnitude is not pinned by the postulates, so the signal’s amplitude is predicted in form (the d^{-6} scaling and the coefficient $\xi = 90/\pi^2$) but not in absolute size.

10.9. Wilson coefficients and electroweak precision: $\Delta S = \Delta T = 0$

PST’s structural commitments fix the dim-6 SMEFT Wilson coefficients that control the electroweak S and T parameters to zero exactly, so PST contributes no new-physics correction to the global electroweak fit (Computation 10, `computation_10.py`, §10).

The argument is structural. PST’s Landau-Ginzburg modal potential $V(H^\dagger H) = -\varepsilon (H^\dagger H) + \frac{1}{4}(H^\dagger H)^2$ is a function of the $SU(2)_L \times SU(2)_R$ singlet $H^\dagger H$ only; it is manifestly custodial-

symmetric. The dim-6 operators $\mathcal{O}_{HWB} = (H^\dagger \sigma^a H) W^{a\mu\nu} B_{\mu\nu}$ and $\mathcal{O}_{HD} = (H^\dagger D_\mu H)^\dagger (H^\dagger D^\mu H)$ are the unique custodial-breaking dim-6 operators that mix the $SU(2)_L$ triplet of W with B and that modify the Higgs kinetic term respectively. Their coefficients c_{HWB} and c_{HD} encode the amount of custodial breaking induced by integrating out heavy modes above the matching scale. Since PST's substrate adds no custodial-breaking new physics above M_* (the LG potential is custodial-symmetric, and the substrate modes integrated out between any μ and M_* are higher harmonics of H itself), the matching-scale Wilson coefficients are $c_{HWB}(M_*) = c_{HD}(M_*) = 0$, exactly.

Consequently the PST contribution to the Peskin–Takeuchi parameters is $\Delta S_{\text{PST}} = (16 \sin \theta_W \cos \theta_W / \alpha_{\text{em}}) \cdot v^2 / M_*^2 \cdot c_{HWB}(M_*) = 0$ and similarly $\Delta T_{\text{PST}} = 0$. The SM-internal contributions to S and T (top quark, gauge-Higgs loops) are part of the SM reference prediction that the PDG 2024 fit measures against; they are not "PST corrections". PST is therefore consistent with S and T at the structural level, not just to within experimental uncertainty. The bounds $|\Delta S| < 0.20$, $|\Delta T| < 0.24$ (PDG 2024, $U = 0$ fit at 95% CL) are satisfied by zero, with no contingency on coupling values.

Other SMEFT operators are pushed to loop level by PST's structural commitments (no new gauge bosons, no new fermions, the substrate modes are higher harmonics of the Higgs). Four-fermion contact operators $(\bar{\psi} \gamma_\mu \psi)(\bar{\psi} \gamma^\mu \psi)$ require new vector or scalar mediators, which PST lacks; their tree-level Wilson coefficients vanish, and the loop-suppressed values give effective scales $\Lambda_{\text{eff}} = M_* / \sqrt{c} \approx 30$ TeV, comfortably above the LHC Drell–Yan high-mass bound of ~ 30 TeV. Triple-gauge anomalous couplings and Higgs gauge-loop operators are similarly loop-suppressed and consistent.

The astrophysical bounds (SN1987A neutron-star cooling, large-extra- dimension graviton emission) do not transfer to PST because PST has no Kaluza–Klein tower: the internal factor F is finite noncommutative with sub-fermi associated length $\ell_* \approx 1.25 \times 10^{-19}$ m, twenty orders of magnitude below the macroscopic radii probed by ADD-type bounds.

The empirical consistency question raised earlier in this section is therefore resolved structurally: $M_* = 4\pi m_h \approx 1.573$ TeV is not in tension with any existing collider or astrophysical bound, by construction.

10.10. Correction budget for the Casimir measurement

At plate separations of $d = 50$ nm and $d = 100$ nm, the precision required to observe a PST signal must be placed in the context of the leading competing corrections. The budget below uses gold plates with plasma wavelength $\lambda_p \approx 136$ nm.

Correction	Relative size at 50 nm	Relative size at 100 nm
Standard Casimir (d^{-4})	1 (reference)	1
Finite conductivity (Lifshitz, $\sim \lambda_p/d$)	$\sim 270\%$	$\sim 136\%$
Surface roughness ($\sim \sigma^2/d^2$ for rms $\sigma \approx 1$ nm)	$\sim 0.04\%$	$\sim 0.01\%$
Thermal/Lifshitz ($k_B T d / \hbar c$, at 300 K)	$\sim 0.5\%$	$\sim 1.0\%$
PST d^{-6} at $d_0 = 4.92$ nm (representative)	$\sim 8.8\%$	$\sim 2.3\%$
PST d^{-6} at $d_0 = 1.1$ nm	$\sim 0.44\%$	$\sim 0.12\%$

The dominant correction at small separations is the finite-conductivity Lifshitz term, which at $d = 50$ nm is *larger* than the standard Casimir force itself; it does not, however, mimic a d^{-6} power law. The Lifshitz formula with the measured optical data of the plate material predicts this correction to better than 1% accuracy [15, 17] and can be subtracted. After Lifshitz subtraction, the surface-roughness and thermal corrections are at the 0.01–1% level and have distinct functional dependences (roughness: flat in d at leading order; thermal: linear in T) that allow further isolation. The PST d^{-6} signal is then the residual. At a representative coherence length $d_0 \approx 4.92$ nm, the correction is $\sim 8.8\%$ at $d = 50$ nm — well above the 1% precision floor of current torsion-pendulum instruments and robustly detectable with next-generation setups.

10.11. Layman’s Summary

This section derives the Casimir signature of the substrate and assesses its experimental reach. Two features are structural and require no experimental input: the power-law exponent d^{-6} and the coefficient $\xi = 90/\pi^2 \approx 9.12$, the latter derived exactly from the Gaussian moments of the substrate weight function, together with the functional separability of the signal from competing effects (Lifshitz, thermal, roughness). The absolute amplitude scales as $(d_0/d)^2$ and is therefore set by the coherence length d_0 of the modal condensate, whose magnitude is a condensate property rather than an output of the postulates; the electroweak modal scale is $M_* = 4\pi m_h \approx 1573$ GeV from m_h alone. For a representative $d_0 \approx 4.92$ nm the predicted correction is $\sim 8.8\%$ at 50 nm and $\sim 55\%$ at 20 nm, well within reach of next-generation precision Casimir experiments.

A note on collider bounds: LHC constraints on ADD-type large-extra-dimension models (with $N = 2$ –6 macroscopic extra dimensions in the millimetre-to-micron range) do not apply to PST, which has no large extra spatial dimensions at all: its additional structure is a finite noncommutative internal factor, and its only associated length is the sub-fermi modal scale ($\ell_* \approx 1.25 \times 10^{-19}$ m), twenty orders of magnitude smaller than the dimensions probed by ADD bounds. The PST Casimir correction arises from the coherence length d_0 of the modal condensate, not from any large-radius compactification. The relevant experimental arena is precision Casimir measurements at nanometre separations, not collider searches for

Kaluza–Klein towers in the ADD regime.

A further structural identification follows from the vacuum topology of Section 8: the emergent spatial geometry is a 3-sphere (S^3), whose isometry group $\text{SO}(4) = \text{SU}(2)_L \times \text{SU}(2)_R$ contains the weak $\text{SU}(2)_L$, with the directed modal threshold selecting the left factor (Section 14.2).

11. Cosmology: Expansion and Acceleration from Ongoing Sublimation

11.1. PST’s reframing of the standard cosmological picture

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 delivers the cosmological-constant sign structurally (Mechanism B of §7.14; Mechanism A gives the standard quantum-vacuum contribution at the SM hierarchy scale).

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

The precausal configuration space is equipped with the Bernoulli product measure $\mu = \bigotimes_{a \in D} \text{Bern}(1/2)$, equation (21), 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 [69], 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 [66]:

$$ds^2 = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right], \quad (154)$$

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.

11.3. 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 [72, 69] 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 $\Omega_k \rightarrow 0$ dynamically throughout the expansion, removing any fine-tuning requirement. No inflaton field or separate scalar sector is needed.

11.4. 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 $\Delta\mathcal{F} = \mathcal{F}_{\text{pre}} - \mathcal{F}|_{\mathcal{V}} > 0$, 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

$$\rho_{\text{inst}} = \frac{j_\tau \Delta\mathcal{F}}{c^2}. \quad (155)$$

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:

$$\dot{\rho}_{\text{inst}} = 0. \quad (156)$$

11.5. 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,

$$\dot{\rho} + 3H(\rho + P) = 0, \quad (157)$$

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

$$P_{\text{inst}} = -\rho_{\text{inst}}, \quad w \equiv \frac{P_{\text{inst}}}{\rho_{\text{inst}}} = -1. \quad (158)$$

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

$$\Lambda_{\text{PST}} = 8\pi G \rho_{\text{inst}} > 0. \quad (159)$$

This is positive, de Sitter, not anti-de Sitter. Section 7.14 identifies this as Mechanism B in the joint structural resolution of the cosmological-constant sign; Mechanism A (quantum vacuum on $\mathcal{V} \cong S^1$) gives a second, SM-inherited positive contribution. The sombrero-valley energy $\mathcal{F}|_{\mathcal{V}} < 0$ is internal to the precausal substrate and does not couple directly to emergent gravity as Λ ; the apparent AdS conclusion of earlier drafts conflated $\mathcal{F}|_{\mathcal{V}}$ with the cosmological constant. 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. The magnitude of Λ_{obs} remains an open question; PST inherits the SM cosmological-constant hierarchy problem but, unlike SM, identifies a structurally PST-specific positive contribution ρ_{inst} that could in principle dominate at small $j_\tau \Delta\mathcal{F}$.

11.6. Friedmann equations and cosmic acceleration

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

FRW metric (154) yields the standard Friedmann equations:

$$H^2 = \frac{8\pi G}{3} (\rho_m + \rho_r + \rho_{\text{inst}}) - \frac{kc^2}{a^2}, \quad (160)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3} (\rho_m + 2\rho_r - 2\rho_{\text{inst}}), \quad (161)$$

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 (161) reduces to

$$\frac{\ddot{a}}{a} \approx \frac{8\pi G}{3} \rho_{\text{inst}} > 0, \quad (162)$$

giving late-time acceleration with no additional assumption. The role of ρ_{inst} in PST is structurally identical to Λ in Λ CDM, but the interpretation differs: in Λ CDM, 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.

11.7. The rate problem

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

$$j_\tau = \int_{\mathcal{S}} \mathbf{1}[T(C) \geq \tau] d\mu(C), \quad (163)$$

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 rate-integral computation remains open. The observed value $\Lambda_{\text{obs}} \approx 1.1 \times 10^{-52} \text{ m}^{-2}$ provides a target: $j_\tau \cdot \Delta\mathcal{F} = \Lambda_{\text{obs}} c^4 / 8\pi G \approx 5.4 \times 10^{-10} \text{ J m}^{-3}$. A successful computation of j_τ would determine Λ_{PST} from first principles and, together with the Casimir prediction, constitute a second independent test of the modal-condensate structure of Section 10.

11.8. Observational commitments

Beyond the standard Λ CDM 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 [69] and distinguishable from dynamical dark energy ($w \neq -1$) by forthcoming surveys (Euclid, DESI).
4. **Spatial flatness.** The driving of $\Omega_k \rightarrow 0$ by ρ_{inst} and the absence of initial fine-tuning predict $\Omega_k = 0$ to high precision, consistent with current CMB bounds [69].
5. **Cosmic acceleration.** Equation (162) gives late-time acceleration consistent with Type Ia supernova data [67, 68] without introducing new fields beyond those required by the precausal structure.

11.9. Comparison with steady-state cosmology

The continuous creation of energy density in Bondi–Gold–Hoyle steady-state cosmology [70, 71] 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 [72] 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 [69] and the transfer-function structure of Λ CDM.

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.

11.10. 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

12.1. PST as foundation, not competitor

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.

12.2. 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 [45]. 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.

12.3. 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 remaining open problem in the QFT sector is the Born rule, which has not yet been derived from the gradient structure on configuration space. Spin-statistics is derived structurally in Section 8 from the CAR algebra on the Boolean distinction primitive, and the interacting vacuum is treated in Section 7.20.

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.

12.4. Thermodynamics and Black Hole Entropy

The thermodynamics of black holes, initiated by Bekenstein [46] and Hawking [47], 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 [19] 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 [18] 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.

12.5. 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.

12.6. 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.

12.7. 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 one vacuum is occupied 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 both arrive at the number ten, but for unrelated reasons, and this coincidence should not be mistaken for kinship. String theory's ten is the dimension of a smooth target manifold, fixed by cancellation of the worldsheet conformal anomaly, with six dimensions compactified on a Calabi-Yau manifold. PST's ten is the total KO-dimension of a real spectral triple, a KO-theoretic invariant defined modulo eight by Bott periodicity, the sum of four spacetime dimensions and a KO-dimension-six *finite noncommutative* internal factor with no spatial extent. PST is therefore not higher-dimensional in the spatial sense at all: there are no compactified extra spatial dimensions, no Calabi-Yau, no worldsheet, and no string landscape. The internal factor is the finite algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ carrying the gauge and matter structure, not a hidden geometry. Where string theory imposes ten from the outside as a consistency requirement and then hides six dimensions, PST derives ten as the unique minimum total KO-dimension compatible with a Lorentzian fermion action, given the substrate's KO-dimension six and four-dimensional spacetime. PST has no landscape: the tension configuration C that produces this universe is a specific point in the modal potential, not a choice from an exponentially large set. The question of why strings is not asked in PST, because strings are not primitive and the fundamental object is the distinction substrate.

12.8. 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.

12.9. Large Extra Dimensions and Braneworld Models

PST has no extra *spatial* dimensions at all: its internal factor is finite and noncommutative, of KO-dimension six, with no spatial extent, and the scale d_0 is the Landau-Ginzburg vacuum (Casimir) coherence scale, not a compactification radius. The comparison with ADD below is therefore *a fortiori*: collider and astrophysical bounds on large extra spatial dimensions do not constrain a theory whose internal structure is a finite noncommutative algebra rather than a hidden geometry. The substantive empirical content, the d^{-6} Casimir correction with parameter-free exponent and coefficient $\xi = 90/\pi^2$, is structural; its amplitude is set by d_0 , a property of the modal condensate rather than any compactification.

The Arkani-Hamed-Dimopoulos-Dvali (ADD) model [16] 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 [30] 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.

12.10. Causal Dynamical Triangulations

Causal Dynamical Triangulations (CDT) [31] 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 flow 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.

12.11. The Holographic Principle and AdS/CFT

The holographic principle, proposed by 't Hooft [32] 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 [33] 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 $\rho : D \rightarrow M$ 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.

12.12. Non-Commutative Geometry

Non-commutative geometry, developed by Connes [48], 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.

12.13. 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 this 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 this 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 (134): 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.

12.14. Spencer-Brown and Distinction-Based Ontologies

G. Spencer-Brown's *Laws of Form* [59] 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* [65] 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 pre-causal 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.

12.15. Wheeler’s “It from Bit” and the Mathematical Universe

John Wheeler’s “it from bit” thesis [60] 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 pre-causal configuration that obtain prior to any apparatus.

Tegmark’s Mathematical Universe Hypothesis [61] carries mathematical realism further, proposing that all self-consistent mathematical structures are physically real and that this 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.

12.16. Structural Realism and Bohm's Implicate Order

Ontic structural realism (OSR), as developed by Ladyman [62] and French [63], 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 [64] 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 $\Phi : \mathbf{S} \rightarrow \mathbf{G}$, which projects the holistic pre-causal 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 pre-causal substrate is the implicate order, and Φ is the explicate projection, with the bifurcation condition $T(C) > \tau$ specifying precisely when projection occurs.

12.17. Layman's Summary

This section places PST in the landscape of existing physical and philosophical theories. General Relativity is recovered in full, given the three independent postulates of Section 1.7 (P4 is derived): the field equations, the equivalence principle, and the Lorentzian signature follow as theorems rather than additional 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

13.1. The shape of the argument

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.

13.2. 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 structural form of physical reality (spacetime, causality, matter, energy, gravity, quantum behaviour, angular momentum, and orbital mechanics) follows by logical entailment, with the macroscopic dimension $n = 4$ enforced by the Connes-spectral-triple framework assumption and Lorentzian signature $(-, +, +, +)$ following as a theorem from hyperbolic well-posedness of the Goldstone wave equation (Computation 2 of Section 14.2). Each downstream concept emerges necessarily from the preceding ones; the chain introduces no further free choice. It 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 this universe. These depend on the particular tension configuration C that produced this universe and require knowledge of C beyond its structural form. What PST derives is the framework within which all such specific facts are intelligible.

The central claim has a clearly delimited scope. The genuinely open item collected in Section 14.2 is the three-generation count (item C.2); the Dirac-convergence step (item A.ii) is numerically closed under the relaxed norm-equivalence criterion of equation (55) via the refined symmetric-

monomial bridge of Computations 39–41. The Yukawa/CKM contingency (§8.18, a no-go theorem placing those observables in the contingent $T(C)$ rather than the structural layer), the chiral-bimodule selection (§8.19), and the empirical-consistency calculation for $M_* \approx 1.573$ TeV (§10.9) were previously open and are now closed; they fall inside what the postulates derive. The headline result is that (D, δ) plus the LG potential, together with the framework assumption that their product with the Mosco-limit spacetime is a Connes real spectral triple, derive everything downstream by logical entailment, with the open work isolated to one named place.

13.3. 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 $\Phi : \mathbf{S} \rightarrow \mathbf{G}$.

The functor Φ produces a Lorentzian manifold (M, g) . The signature $(-, +, +, +)$ follows by hyperbolic well-posedness of the Goldstone wave equation (Computation 2, Section 14.2): Lorentzian is the unique signature admitting a well-posed Cauchy problem and a conserved Noether current. The remaining analytic open item is the Dirac-convergence step to the unique spacetime Dirac operator (Section 14.2). 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 (27). The tension-to-metric map is equation (29). The stress-energy tensor is equation (35). The conservation law is equation (36). The Noether charge is equation (66). The chain is not a qualitative narrative but a sequence of mathematical constructions, each of which follows from the previous given the three independent postulates of Section 1.7 (with P4 derived).

13.4. 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 [41] 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, $\text{causality} = f(g) = f(\Phi(C))$; 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.

13.5. 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.

13.6. 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 (27); the functor Φ is constructed in three explicit steps via the realization map ρ and the tension-to-metric assignment of equation (29); the projection operator Π is equation (30); the field equations follow as equation (36); the vacuum manifold radius is $r_0 = \sqrt{\varepsilon/(2b)}$; the Noether charge is equation (66). 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 (140). Section 10 identifies d_0 with the Landau-Ginzburg coherence length of the modal condensate; the electroweak modal scale is fixed by the Higgs mass as $M_* = 4\pi m_h \approx 1.573$ TeV. The correction at $d = 50$ nm is $\sim 0.97\% |\xi|$ for a coherence

length of a few nanometres; with $\xi \in [1, 10]$, this falls at or above the 1% level accessible to next-generation precision Casimir experiments. Existing measurements at 160–750 nm constrain $d_0 < 16$ nm. A ratio test between measurements at two separations, equation (144), provides a clean, apparatus-independent signature distinguishable from finite-conductivity, surface roughness, and thermal corrections. The power-law exponent -6 is a parameter-free structural consequence; observing a d^{-6} correction determines the amplitude and hence d_0 , while a null result at 1% level at 50 nm excludes $\xi > 1$ and constrains the remaining coefficient.

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 *qualitatively consistent* with the independent result of Causal Dynamical Triangulations [31]. PST and CDT both report a flow from a higher spectral dimension at short scales to four at large scales; whether the specific scale and slope agree quantitatively has not been checked and the qualitative agreement should not be read as a numerical corroboration.

13.7. 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. The gauge group $SU(3) \times SU(2) \times U(1)$ arises as the unimodular unitaries of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ (Section 14.2); what remains open in the matter sector is the three-generation count (item C.2). 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 electroweak modal scale is fixed unconditionally as $M_* = 4\pi m_h \approx 1.573$ TeV; the coherence length d_0 is a condensate property whose magnitude is not pinned by the postulates alone, while the dimensionless coefficient ξ is bounded to $[1, 10]$ from the Gaussian approximation to the modal kernel, with the exact value requiring the non-Gaussian Boolean lattice corrections to K . 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.

13.8. 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.

13.9. Layman's Summary

Precausal Substrate Theory derives spacetime, causality, matter, the four fundamental forces, and the statistical uniformity of the cosmos from three independent postulates: the distinction primitive, asymmetric tension, and the Landau-Ginzburg modal potential. The Laplace-type projection (P4) is derived from these three via Mosco convergence of Boolean Dirichlet forms. All structure beyond those three commitments follows as a consequence; 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. Summary of Results

14.1. Research problems resolved

This section summarises the research problems the theory resolves, grouped by sector. Open work follows in §14.2. Each result below is established in the body and supported by an explicit verification.

The foundational object. The foundational locus of PST is a real spectral triple of total KO-dimension $4 + 6 = 10 \equiv 2 \pmod{8}$, the value compatible with a Lorentzian fermion action. The substrate's real structure has KO-dimension six conditional on the parity-selected sub-limit (an odd number of distinction bits; Computation 3). The Mosco limit of the Boolean Dirichlet form gives a four-dimensional Lorentzian spacetime of KO-dimension four (Computation 4); their product is the foundational triple of total KO-dimension ten. The four and the six are the minimal split compatible with the verified spectral-triple structure and the parity selection. The ten is a KO-theoretic invariant $\pmod{8}$, not a count of spatial dimensions and not the critical dimension of superstring theory. The construction is verified explicitly (`computation_03.py`): the product $M \times F$ is a genuine real spectral triple with the five pieces (A, H, D, J, γ) written out and the axioms checked.

Gravity. The Einstein-Hilbert action is the leading term of the spectral action $\text{Tr } f(D/\Lambda)$, with Newton's constant entering as a consistency relation among (m_h, M_P, v) rather than a free input (`computation_05.py`). The spin structure on $\mathbb{R} \times S^3$ is unique, so the spacetime Dirac operator is canonical and the one analytic residual is a convergence theorem (`computation_04.py`).

Quantum mechanics. The canonical commutation relations follow from the gradient structure on configuration space, via Mosco convergence of the Boolean Dirichlet forms. The differentiable structure is supplied by the Boolean gradient derived from the symmetric-difference metric, and the Bernoulli measure $\mu = \bigotimes_{a \in D} \text{Bern}(1/2)$ is uniquely fixed by relabelling and complementation invariance, with no additional primitive.

The Standard Model. The gauge group $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$ is the group of unimodular unitaries of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, with the gauge bosons and Higgs as inner fluctuations of the Dirac operator (`computation_06.py`). The colour $\text{SU}(3)_c$ is the $M_3(\mathbb{C})$ factor; one generation's colour and charge assignments follow from the octonionic origin of A_F . The chiral $\text{SU}(2)_L$ is conditional on a chiral bimodule for A_F : given it, the order-one condition selects the Yukawa D_F uniquely (`computation_07.py`); the directed modal threshold is the live candidate to supply it (`computation_08.py`), with the bimodule-support residual in Section 14.2.

The Dirac-convergence step. The lift from the substrate's noncommutative-geometric reading to its $\mathbb{R} \times S^3$ Mosco limit is closed in both sub-parts (§7.7). The section round-trip lemma – a quantum Gromov–Hausdorff statement analogous to [100] Lemma 3.2 – closes under the Fréchet smooth Lip-norm $L_F = \sup_j \|\text{ad}_D^j(\cdot)\|/j!$ at exponential rate $\gamma_D \approx 0.37 e^{-0.28D}$

(Computations 20–23), giving QGH convergence of the Walsh substrate algebras to $C(S^3)$ via [97] Theorem 5.2. The *single-mode* L_{round} closure is now a theorem (Computation 50): the Jordan–Wigner Clifford structure gives $D_{\text{sub}}^2 = D \cdot I$, hence $\|\text{ad}_D^j(\chi_S)\|_{\text{op}} \leq (2\sqrt{D})^j$ uniformly, and by Stirling $\|\chi_S\|_{\text{op}}/L_F(\chi_S) \leq e^{-2\sqrt{D}} \rightarrow 0$ rigorously. For the constant-coefficient tail $T_{\text{tail}} = \sum_{|S|>k_D} \chi_S$, Computation 51 establishes the closed-form block-decomposition identity $\|\text{ad}_D^j(T)\|_{\text{op}} = (2\sqrt{D})^j \cdot \max(\|T_{+-}\|, \|T_{-+}\|)$ for every T and every $j \geq 1$, reducing the full L_{round} closure to an explicit operator-norm bound on the off-diagonal block $P_+T_{\text{tail}}P_-$ in the D_{sub} eigenbasis. The true asymptotic rate is $\gamma_D = O(D^{1/4}e^{-2\sqrt{D}})$, with the empirical $0.37e^{-0.28D}$ a pre-asymptotic fit. The Dirac-aware lift closes under the relaxed norm-equivalence criterion of equation (55) via the refined symmetric-monomial bridge $b_C(\chi_S) = T_{f_{|S|}}$ with $f_k = (z_0z_1)^{m(k)-1} + \beta(k)(z_0z_1)^{m(k)} + \alpha(k)(z_0z_1)^{m(k)+1}$ (Computations 39–42); the bare $m = 1$ monomial gives ratio = 2 = $2\sqrt{1}$ exactly at $k = 1$, two-monomial superposition closes even k with gap below 10^{-3} , and three-monomial superposition closes odd k and extends uniformly across $k = 1, 2, \dots, 12$ with maximum gap 3.6×10^{-4} at $k = 9$ (closure exact to grid precision at $k = 1, 3, 5, 12$). The construction gives $\gamma(k)$ that is D -independent. The relaxed criterion is strictly sufficient for the spectral-action and gauge-group recoveries in the body, since each invokes spectral data of D at finite D rather than full operator equality of bridge commutators (§7.7, *Why the relaxed criterion suffices for PST*).

Octonionic-SU(3) closure. The octonionic gauge structure is now closed at the rigorous level: §8.6 gives the Casimir-discriminated SU(3) decomposition of $M_8(\mathbb{C}) = 6 \cdot \mathbf{1} + 5 \cdot \mathbf{3} + 5 \cdot \bar{\mathbf{3}} + 2 \cdot \mathbf{8} + \mathbf{6} + \bar{\mathbf{6}}$ to machine precision (Computation 17); the embedding choice is justified by Spin(6)-invariance (§8.8, Computation 18): the volume element $\omega = L_{e_1} \dots L_{e_6}$ is the unique (up to scalar) Spin(6)-invariant element of positive grade in $\text{Cl}(0, 6) \otimes \mathbb{C}$, so the Furey primitive idempotent $f = (I + i\omega)/2$ is the unique Spin(6)-invariant primitive idempotent, fixing the embedding up to chiral swap. Conjecture 1 (the algebra identity) is closed via Computations 3, 21, and Furey 2014 (see also the **Conjecture 1 closure** block below). The generation-count derivation $N_{\text{gen}} = 3$ is closed via the V_7 -valued directional structure of P1–P3 and the $\hat{\tau}$ -selection at modal sublimation (§8.17): the three quaternionic sub-algebras of \mathbb{O} containing the $\hat{\tau}$ -direction are identified with the three generations of matter content (Computations 46–48).

Matter and the empirical tail. The Boolean $\{0, 1\}$ distinctions are fermionic occupation numbers, making spin-statistics structural. The instantiated vacuum is a 6-sphere of degenerate directions in V_7 ; its 2-plane projection is the familiar ring of minima, so orbital motion is the only stable ground state. The falsifiable d^{-6} Casimir correction has a parameter-free power-law exponent and coefficient $\xi = 90/\pi^2$ at the substrate coherence scale d_0 ; the electroweak modal scale is $M_* = 4\pi m_h \approx 1.573$ TeV.

Conjecture 1 (the algebra identity), CLOSED. *P1–P3 imply that the substrate’s real spectral triple has internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ at KO-dimension six.*

This statement is closed in §8.8. The closure chain has three computationally verified pieces:

- (1) *Computation 3* (`computation_03.py`): on the parity-selected odd- D sub-limit, the substrate's complementation-and-parity real structure delivers $\gamma_K = Z^{\otimes D}$ with $\gamma_K^2 = +I$, sign triple $(\varepsilon, \varepsilon', \varepsilon'') = (+, +, -)$, the Connes Lorentzian-compatible KO-6 value.
- (2) *Computation 21 §6* (`computation_19.py`): in the Jordan-Wigner Majorana representation, $\gamma_K = i\omega_{JW}$ exactly ($\|\gamma_K - i\omega_{JW}\| = 0$ to machine precision, where ω_{JW} is the volume element of the $\text{Cl}(0,6)$ Majorana algebra). The substrate's parity grading IS the $\text{Cl}(0,6)$ chirality grading, and the chirality projector $f = (I + i\omega)/2$ is delivered directly, this is exactly Furey's primitive idempotent.
- (3) *Furey 2014* [102]: the commutant of the matter representation on the minimal left ideal $\text{Cl}(0,6) \cdot f$ is $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$.

Supporting verification: Computation 16 established $\overleftarrow{\mathbb{O}} \cong \text{Cl}(0,6) \otimes \mathbb{C} \cong M_8(\mathbb{C})$, and Computation 18 certified that f is the unique $\text{Spin}(6)$ -invariant primitive idempotent (singlet count $(1, 0, 0, 0, 0, 0, 1)$ over grades 0–6 in the $\text{Spin}(6)$ decomposition of $\text{Cl}(0,6)$); Computation 19 §4 confirmed the Hermitian $\text{Spin}(6)$ -invariant operators on H_F form exactly $\text{span}\{I, \gamma\}$ by Schur's lemma on the irreducible chiral halves $H_+ \cong \mathbf{4}$, $H_- \cong \bar{\mathbf{4}}$. These are not load-bearing for the closure but provide independent certification that f is the canonical structural choice.

The derivation makes no appeal to the Chamseddine-Connes-Marcolli classification. Conjecture 1 is therefore closed for both readings of the abstract: under reading (i) (grounding) the algebra is inherited; under reading (ii) (derivation from P1–P3) the algebra is derived via the chain above.

14.2. Open research work

All structural results derived from the foundational object are collected in §14.1, distributed across the body of the paper at the section references listed there. Two precise analytical items remain open in the rigorous-proof tail of the Dirac-convergence-step closure (§7.7); the single-mode L_{round} closure (Computation 50) and the block-decomposition reduction (Computation 51) are already rigorous theorems and are accounted in §14.1.

Off-diagonal block bound for the constant-coefficient Walsh tail. The block-decomposition identity of Computation 51 reduces the full L_{round} closure to the explicit operator-norm bound $\max(\|(T_{\text{tail}})_{+-}\|_{\text{op}}, \|(T_{\text{tail}})_{-+}\|_{\text{op}}) \leq C \cdot 2^{D-1}$ for an explicit constant $C > 0$. Numerically the bound is $2^{D-1} - D$ exactly across $D \in \{4, 5, 6, 7, 8\}$ at cutoff $k_D = 2$; the rigorous derivation is a Fourier-on-Boolean-cube estimate on the explicit operator $P_+ T_{\text{tail}} P_- = \frac{1}{4}(I + D_{\text{sub}}/\sqrt{D})T_{\text{tail}}(I - D_{\text{sub}}/\sqrt{D})$.

L_{comm} refined-bridge rigorous closure. The refined symmetric-monomial bridge (Computations 39–42) closes the relaxed norm-equivalence criterion to gap $< 10^{-4}$ across $k = 1, \dots, 12$ uniformly in D . The remaining work is the rigorous proof: Berezin-transform compatibility of the bridge with the weighted-Bergman spectral-triple structure ([100]), spectral-action invariance under the refined bridge, and the uniform bound across $k \leq k_D$ for all D . Estimated

effort: research-paper-level operator-algebras analysis (1–2 months).

14.3. Layman’s Summary

PST starts from one primitive (property differentiation) with directional content in a 7-dimensional algebra V_7 fixed by Hurwitz’s classification. From the three postulates the substrate is a real spectral triple whose Mosco limit pairs four-dimensional spacetime $\mathbb{R} \times S^3$ with an internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ to total KO-dimension ten, the Connes Standard-Model value.

Out of this single object the paper derives: the Lorentzian signature of spacetime; the canonical commutation relations of quantum mechanics; the Einstein–Hilbert action from the spectral action; the Standard Model gauge group as the unimodular unitaries of A_F , with the Higgs and gauge bosons as inner fluctuations of the Dirac operator; one generation of fermion content from the octonionic minimal ideal; and three generations from the $\hat{\tau}$ -selection of a unit direction in V_7 (the three quaternionic sub-algebras of \mathbb{O} containing $\hat{\tau}$). The Yukawa hierarchies, CKM mixing, the cosmological constant magnitude, and the substrate coherence scale d_0 remain contingent.

Two analytical items remain open in the rigorous-proof tail of the Dirac-convergence step (§14.2): an explicit operator-norm bound on the off-diagonal block of the Walsh tail and the Berezin-transform compatibility of the refined symmetric-monomial bridge. The structural single-mode L-round closure and the block-decomposition reduction are already rigorous theorems (Computations 50, 51).

The Dirac-convergence step is numerically closed in two parts. The quantum Gromov–Hausdorff section round-trip lemma is closed by the Fréchet smooth Lip-norm at exponential rate $\gamma_D \approx 0.37 e^{-0.28D}$ (Computation 23), giving QGH convergence of the Walsh substrate algebras to $C(S^3)$ via [97] Theorem 5.2. The Dirac-aware lift is closed under the relaxed norm-equivalence criterion ((55)) by the refined symmetric-monomial bridge $b_C(\chi_S) = T_{f_{|S|}}$ with $f_k = (z_0 z_1)^{m(k)-1} + \beta(k)(z_0 z_1)^{m(k)} + \alpha(k)(z_0 z_1)^{m(k)+1}$ (Computations 39–41). The refined bridge gives $\gamma(k)$ D -independent and below 10^{-4} at confirmed k , eliminating the k_D cutoff-transition discontinuity exposed by Computation 38 and the homomorphism-gap obstruction of Computation 35. The relaxed criterion is strictly sufficient for the spectral-action and gauge-group recoveries in the body, since each invokes spectral data of D at finite D rather than full operator equality of bridge commutators. The remaining analytical work for the Dirac-convergence step is the rigorous proof of the numerical closure: Berezin-transform compatibility of b_C with the weighted-Bergman spectral-triple structure ([100]), spectral-action invariance under b_C , and the uniform bound across $k \leq k_D$ for all D .

The three-generation count: the structural derivation of $N_{\text{gen}} = 3$ generations on top of the verified octonionic-SU(3) gauge structure (§8.6); two candidate routes (Furey 2018, or Computation 1 contingency on the Bernoulli substrate) are consistent with the framework, neither obstructing it. All other items previously open (the chiral bimodule, the Yukawa/CKM

contingency, and the empirical M_* bound) are closed and documented in the body subsections §8.19, §8.18, and §10.9 respectively.

Appendix A: Notation Conventions

The following symbols recur throughout the paper. Section references point to where each is first introduced.

Substrate and primitives.

Symbol	Meaning	Section
(D, δ)	Discrete distinction substrate: a set D with the distinction primitive δ	§2
$\mathcal{P}(D)$	Boolean configuration space of subsets of D	§2
C	A precausal configuration, $C \subseteq D$	§2
$T(C)$	Asymmetric-tension functional on $\mathcal{P}(D)$	§2
μ	Bernoulli product measure on $\mathcal{P}(D)$, $\bigotimes_a \text{Bern}(1/2)$	§3

Modal threshold and order parameter.

Symbol	Meaning	Section
τ	Modal threshold value of the tension	§3
$\varepsilon = T(C) - \tau$	Excess tension (the control parameter)	§3
ψ	Modal order parameter (complex scalar)	§3
$\psi^* = \pm\sqrt{2\varepsilon}$	Order parameter at the potential minimum	§3
$F[\psi, \varepsilon]$	Modal potential (Landau-Ginzburg form)	§3
\mathcal{V}	Vacuum manifold; for $\varepsilon > 0$, $\mathcal{V} \cong S^6 \subset V_7$ (a 6-sphere of directions); the 2-plane projection is the familiar ring of minima	§3
r_0	Vacuum radius, $r_0^2 = \varepsilon/(2b)$	§3

Modal sublimation and projection.

Symbol	Meaning	Section
$\Phi : \mathcal{S} \rightarrow \mathcal{G}$	Modal-sublimation functor (precausal \rightarrow geometric)	§4
ρ	Realisation map from substrate to spacetime	§4
Π	Projection operator, substrate \rightarrow four-dimensional field theory	§5
κ_{proj}	Projection coefficient (Bernoulli-measure integral)	§7

Foundational object: the real spectral triple.

Symbol	Meaning	Section
$M = \mathbb{R} \times S^3$	Four-dimensional Lorentzian spacetime factor (KO-dimension 4)	§1.4
F	Finite noncommutative internal factor (KO-dimension 6)	§1.4
$M \times F$	Foundational real spectral triple (total KO-dimension $\equiv 2 \pmod{8}$)	§1.4
(A, H, D, J, γ)	Spectral-triple data: algebra, Hilbert space, Dirac operator, real structure, chirality grading	§1.4
$A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$	Internal algebra of F	§1.4
$\text{Im}(\mathbb{O}) \cong \mathbb{R}^7$	Imaginary octonions: the algebraic space on which F is built	§8
$\text{Cl}(0, 6)$	Clifford algebra on $\text{Im}(\mathbb{O})$	§8

Standard Model and gauge sector.

Symbol	Meaning	Section
$\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$	Standard Model gauge group; unimodular unitaries of A_F	§8
$N_c = 3$	Colour multiplicity; the rank of $M_3(\mathbb{C})$	§8
N_{gen}	Generation count (open; not derived)	§8
D_F	Yukawa / mass operator on H_F	§8
Y	$\text{U}(1)_Y$ hypercharge	§8
$Z^2 \approx e^{-1}$	Field-normalisation ratio (one-loop near-coincidence)	§8

Scales and constants.

Symbol	Meaning	Section
m_h	Higgs mass (≈ 125.20 GeV)	§7
v	Higgs vacuum expectation value (≈ 246 GeV)	§7
$M_* = 4\pi m_h$	Electroweak modal scale (≈ 1.573 TeV)	§7
$\ell_* = \hbar/M_*$	Modal length scale ($\approx 1.25 \times 10^{-19}$ m)	§10
d_0	Substrate coherence scale (Landau-Ginzburg coherence length)	§10
ℓ_P	Planck length	§10
M_P	Planck mass	§10
G	Newton's gravitational constant (consistency relation in PST)	§7
$\alpha_G \equiv G/d_0^2$	Gravitational-coupling coefficient (in $\hbar = c = 1$)	§7
$\xi = 90/\pi^2$	Casimir-correction coefficient (parameter-free)	§10
Λ	Cosmological constant	§11

Appendix B: Bibliographic 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. M. Lifshitz, “The theory of molecular attractive forces between solids,” *Sov. Phys. JETP* **2**, 73 (1956).
- [18] E. P. Verlinde, “On the origin of gravity and the laws of Newton,” *JHEP* **04** (2011) 029. arXiv:1001.0785
- [19] T. Jacobson, “Thermodynamics of spacetime: The Einstein equation of state,” *Phys. Rev. Lett.* **75**, 1260 (1995). arXiv:gr-qc/9504004
- [20] A. Einstein, “Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie,” *Sitzungsber. Preuss. Akad. Wiss.* **142** (1917). ADS:1917SPAW.....142E
- [21] E. Noether, “Invariante Variationsprobleme,” *Nachr. D. König. Gesellsch. D. Wiss. Zu Göttingen, Math-phys. Klasse*, 235 (1918). arXiv:physics/0503066
- [22] D. Lovelock, “The Einstein tensor and its generalizations,” *J. Math. Phys.* **12**, 498–501 (1971). doi:10.1063/1.1665613
- [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] 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
- [31] J. Ambjørn, J. Jurkiewicz, R. Loll, “Reconstructing the universe,” *Physical Review D* **72** (2005) 064014. doi:10.1103/PhysRevD.72.064014
- [32] G. ’t Hooft, “Dimensional reduction in quantum gravity,” arXiv:gr-qc/9310026 (1993). arXiv:gr-qc/9310026
- [33] 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

- [34] 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.]
- [35] 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
- [36] 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
- [37] R. Geroch, “Domain of dependence,” *Journal of Mathematical Physics* **11** (1970) 437–449. doi:10.1063/1.1665157
- [38] R. M. Wald, *General Relativity*. University of Chicago Press, Chicago, 1984. University of Chicago Press
- [39] G. Ryle, *The Concept of Mind*. Hutchinson, London, 1949. [archive scan]
- [40] H. Putnam, *Reason, Truth and History*. Cambridge University Press, Cambridge, 1981. doi:10.1017/CBO9780511625398
- [41] 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]
- [42] A. Vilenkin, “Creation of universes from nothing,” *Physics Letters B* **117** (1982) 25–28. doi:10.1016/0370-2693(82)90866-8
- [43] J. Goldstone, “Field Theories with Superconductor Solutions,” *Il Nuovo Cimento* **19** (1961) 154–164. doi:10.1007/BF02812722
- [44] P. A. M. Dirac, “The cosmological constants,” *Nature* **139** (1937) 323. doi:10.1038/139323a0
- [45] 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
- [46] J. D. Bekenstein, “Black holes and entropy,” *Physical Review D* **7** (1973) 2333–2346. doi:10.1103/PhysRevD.7.2333
- [47] S. W. Hawking, “Particle creation by black holes,” *Communications in Mathematical Physics* **43** (1975) 199–220. doi:10.1007/BF02345020
- [48] A. Connes, *Noncommutative Geometry*. Academic Press, San Diego, 1994. [free PDF, author’s website]
- [49] A. H. Chamseddine and A. Connes, “The spectral action principle,” *Commun. Math. Phys.* **186** (1997) 731–750, arXiv:hep-th/9606001. doi:10.1007/s002200050126

- [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] A. M. Gleason, “Measures on the closed subspaces of a Hilbert space,” *J. Math. Mech.* **6** (1957) 885–893. doi:10.1512/iumj.1957.6.56050
- [53] B. Efron and C. Stein, “The jackknife estimate of variance,” *Ann. Statist.* **9** (1981) 586–596. doi:10.1214/aos/1176345462
- [54] R. O’Donnell, *Analysis of Boolean Functions*, Cambridge University Press (2014). ISBN 978-1107038325. doi:10.1017/CBO9781139814782
- [55] B. S. DeWitt, *Dynamical Theory of Groups and Fields*, Gordon and Breach, New York (1965).
- [56] P. B. Gilkey, “The spectral geometry of a Riemannian manifold,” *J. Differential Geometry* **10** (1975) 601–618. doi:10.4310/jdg/1214433164
- [57] D. V. Vassilevich, “Heat kernel expansion: User’s manual,” *Phys. Rept.* **388** (2003) 279–360. arXiv:hep-th/0306138
- [58] B. S. DeWitt, “Quantum Theory of Gravity. I The Canonical Theory,” *Physical Review* **160** (1967) 1113–1148. doi:10.1103/PhysRev.160.1113
- [59] G. Spencer-Brown, *Laws of Form*. Allen & Unwin, London, 1969. [archive.org scan]
- [60] 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
- [61] M. Tegmark, “The Mathematical Universe,” *Foundations of Physics* **38** (2008) 101–150. doi:10.1007/s10701-007-9186-9 [arXiv:0704.0646]
- [62] 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
- [63] S. French, *The Structure of the World: Metaphysics and Representation*. Oxford University Press, Oxford, 2014. ISBN 978-0-19-968484-7.
- [64] D. Bohm, *Wholeness and the Implicate Order*. Routledge & Kegan Paul, London, 1980. [archive.org scan]
- [65] 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]
- [66] A. Friedmann, “Über die Krümmung des Raumes,” *Zeitschrift für Physik* **10** (1922) 377–386. doi:10.1007/BF01332580

- [67] 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]
- [68] 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]
- [69] 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]
- [70] 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
- [71] 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
- [72] 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
- [73] P. W. Higgs, “Broken symmetries and the masses of gauge bosons,” *Phys. Rev. Lett.* **13** (1964) 508–509. doi:10.1103/PhysRevLett.13.508
- [74] F. Englert and R. Brout, “Broken symmetry and the mass of gauge vector mesons,” *Phys. Rev. Lett.* **13** (1964) 321–323. doi:10.1103/PhysRevLett.13.321
- [75] S. Weinberg, “A model of leptons,” *Phys. Rev. Lett.* **19** (1967) 1264–1266. doi:10.1103/PhysRevLett.19.1264
- [76] A. Salam, “Weak and electromagnetic interactions,” in *Proceedings of the 8th Nobel Symposium*, ed. N. Svartholm, Almqvist & Wiksell, Stockholm (1968), pp. 367–377.
- [77] S. L. Glashow, “Partial symmetries of weak interactions,” *Nucl. Phys.* **22** (1961) 579–588. doi:10.1016/0029-5582(61)90469-2
- [78] H. Fritzsch, M. Gell-Mann, and H. Leutwyler, “Advantages of the color octet gluon picture,” *Phys. Lett. B* **47** (1973) 365–368. doi:10.1016/0370-2693(73)90625-4
- [79] D. J. Gross and F. Wilczek, “Ultraviolet behavior of non-Abelian gauge theories,” *Phys. Rev. Lett.* **30** (1973) 1343–1346. doi:10.1103/PhysRevLett.30.1343
- [80] H. D. Politzer, “Reliable perturbative results for strong interactions?” *Phys. Rev. Lett.* **30** (1973) 1346–1349. doi:10.1103/PhysRevLett.30.1346
- [81] T. D. Lee and C. N. Yang, “Question of parity conservation in weak interactions,” *Phys. Rev.* **104** (1956) 254–258. doi:10.1103/PhysRev.104.254
- [82] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, “Experimental test of parity conservation in beta decay,” *Phys. Rev.* **105** (1957) 1413–1415.

- [83] S. L. Adler, “Axial-vector vertex in spinor electrodynamics,” *Phys. Rev.* **177** (1969) 2426–2438. doi:10.1103/PhysRev.177.2426
- [84] J. S. Bell and R. Jackiw, “A PCAC puzzle: $\pi^0 \rightarrow \gamma\gamma$ in the σ -model,” *Nuovo Cim. A* **60** (1969) 47–61. doi:10.1007/BF02823296
- [85] G. Aad *et al.* (ATLAS Collaboration), “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” *Phys. Lett. B* **716** (2012) 1–29. doi:10.1016/j.physletb.2012.08.020
- [86] S. Chatrchyan *et al.* (CMS Collaboration), “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,” *Phys. Lett. B* **716** (2012) 30–61. doi:10.1016/j.physletb.2012.08.021
- [87] R. L. Workman *et al.* (Particle Data Group), “Review of particle physics,” *Prog. Theor. Exp. Phys.* **2022** (2022) 083C01. doi:10.1093/ptep/ptac097
- [88] M. J. G. Veltman, “The Infrared–Ultraviolet Connection,” *Acta Phys. Polon. B* **12** (1981) 437–457.
- [89] G. ’t Hooft, “Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking,” in *Recent Developments in Gauge Theories*, NATO ASI Series B59, eds. G. ’t Hooft *et al.*, Plenum Press, New York (1980), pp. 135–157.
- [90] J. Polchinski, “Renormalization and effective lagrangians,” *Nucl. Phys. B* **231** (1984) 269–295. doi:10.1016/0550-3213(84)90287-6
- [91] R. Rajaraman, *Solitons and Instantons: An Introduction to Solitons and Instantons in Quantum Field Theory*, North-Holland, Amsterdam (1982). ISBN 978-0-444-86229-7.
- [92] U. Mosco, “Composite media and asymptotic Dirichlet forms,” *J. Funct. Anal.* **123** (1994) 368–421. doi:10.1006/jfan.1994.1093
- [93] M. E. Taylor, *Partial Differential Equations I: Basic Theory*, Springer, New York (1996). doi:10.1007/978-1-4684-9320-7
- [94] T. Friedrich, “Der erste Eigenwert des Dirac-Operators einer kompakten, Riemannschen Mannigfaltigkeit nichtnegativer Skalar­krümmung,” *Math. Nachr.* **97** (1980) 117–146. doi:10.1002/mana.19800970111
- [95] C. Bär, “The Dirac operator on space forms of positive curvature,” *J. Math. Soc. Japan* **48** (1996) 69–83. doi:10.2969/jmsj/04810069
- [96] M. A. Rieffel, “Matrix algebras converge to the sphere for quantum Gromov–Hausdorff distance,” *Mem. Amer. Math. Soc.* **168** (2004) 67–91. arXiv:math/0108005
- [97] M. A. Rieffel, “Gromov–Hausdorff distance for quantum metric spaces,” *Mem. Amer. Math. Soc.* **168** (2004) 1–65. arXiv:math/0011063. *The bridge construction (Definition*

5.1) and the metric-via-bridge theorem (Theorem 5.2) used to lift the section round-trip lemma to quantum Gromov–Hausdorff convergence.

- [98] F. Latrémolière, “The dual-modular Gromov–Hausdorff propinquity and completeness,” *J. Noncommut. Geom.* (2022). arXiv:1811.04534
- [99] F. Latrémolière, “The Gromov–Hausdorff propinquity for metric spectral triples,” *Adv. Math.* **404** (2022) 108393. arXiv:1811.10843
- [100] T. Bhattacharyya and S. Singla, “Sequences of operator algebras converging to odd spheres in the quantum Gromov–Hausdorff distance,” *preprint* (2022). arXiv:2211.04321
- [101] P. Diaconis and M. Shahshahani, “Time to reach stationarity in the Bernoulli–Laplace diffusion model,” *SIAM J. Math. Anal.* **18** (1987) 208–218. doi:10.1137/0518016
- [102] C. Furey, “Three generations, two unbroken gauge symmetries, and one eight-dimensional algebra,” *Phys. Lett. B* **785** (2018) 84–89. arXiv:1810.10518

Appendix C: Computation Scripts

Each result tagged “verified” in §14.1 is checked symbolically or numerically by the Python script listed below. The scripts live in the project’s `computations/` directory as `computation_NN.py` and can be run independently.

Comp.	Result
1	Yukawa hierarchy.
§8.18	Generation-symmetry no-go theorem for a structural Yukawa hierarchy: substrate-level S_3 invariance forces Yukawa eigenvalues to be S_3 -symmetric in leading order, so the observed hierarchy lives in the contingent $T(C)$ rather than the structural layer (occupying $\approx 31\%$ of the $2^D = 64$ bit budget at $D = 6$).
2	Lorentzian signature.
§7.6	Lorentzian signature $(-, +, +, +)$ from Derrick-style hyperbolic well-posedness of the Goldstone wave equation; the macroscopic dimension $n = 4$ is forced by KO additivity given the verified substrate KO-6 (Computation 3) and the total KO-2 of the product.
3	Spectral triple KO-2.
§1.4	The product $M \times F$ is a genuine real spectral triple of KO-dimension 2.
4	Spin structure on $\mathbb{R} \times S^3$. The unique spin structure on $\mathbb{R} \times S^3$.
§7.6	
5	Einstein-Hilbert and Newton’s G.
§7.9	Einstein-Hilbert and Newton’s G as the gravitational term of the spectral action.
6	Gauge group from A_F.
§8	Gauge group from the unimodular unitaries of A_F ; gauge bosons and Higgs as inner fluctuations of D .
7	Order-one condition.
§8.19	The order-one condition selects the Yukawa D_F ; $SU(2)_L$ acts chirally given the chiral bimodule.
8	Chirality bridge.
§8.19	γ_M does not supply chirality; the directed modal threshold is the live candidate.
9	Dirac convergence scoping.
§7.7	Dirac-convergence scoping in spectral propinquity; numerical companion verifies $Cl(0, 2D)$ anticommutation at every D , Walsh-to-spherical multiplicity flow, and exact low-mode commutator-gap on Walsh-weight ≤ 1 modes; residual reduced to one uniform Sobolev-tail estimate (51).

Comp.	Result
10	Empirical M_* bounds.
§10.8	SMEFT operator-by-operator survey of $M_* = 4\pi m_h$ at PST's tree level vs LHC/EWPO bounds; $\Delta S_{\text{PST}} = \Delta T_{\text{PST}} = 0$ structurally.
11	Chiral bimodule support.
§8.19	Both substrate $\text{su}(2)$ factors have clean block support; threshold selects between parity-conjugate choices.
12	Octonionic-SU(3) consistency check. Compatible with Computations 4, 5,
§8.7	6, 7, 8, 9, 10, 11 under the octonionic embedding.
13	Canonical vs octonionic SU(3). Side-by-side comparison under Computa-
§8.7	tions 1 (§7) and 7.
14	S_3 invariance under octonionic embedding.
§8.7	Computation 1 (§7) under octonionic embedding: individual fluctuations can mix generations but substrate S_3 invariance constrains physical observables; Yukawa contingency preserved.
15	Order-one under octonionic embedding.
§8.7	Computation 7 under octonionic embedding: order-one selecting condition is unitarily covariant; residue $\sim 2 \times 10^{-5}$ (Yukawa) vs ~ 10 (generic) under both embeddings.
16	Furey foundation.
§8.5	Octonion algebra, L_{e_a} Clifford structure, $\overleftarrow{\mathbb{O}} = 64$ -complex-dim verified, and octonionic SU(3) as G_2 -stabiliser of e_7 constructed to machine precision.
17	SU(3) decomposition of $M_8(\mathbb{C})$. Definitive Casimir-discriminated decomposi-
§8.6	tion: $6 \cdot \mathbf{1} + 5 \cdot \mathbf{3} + 5 \cdot \bar{\mathbf{3}} + 2 \cdot \mathbf{8} + \mathbf{6} + \bar{\mathbf{6}}$ matched to Clebsch–Gordan to machine precision.
18	ω as the unique Spin(6) singlet.
§8.8	Spin(6)-invariance in $\text{Cl}(0, 6) \otimes \mathbb{C}$: $\omega = L_{e_1} \dots L_{e_6}$ is the unique (up to scalar) Spin(6)-invariant element of positive grade. Hence $f = (I + i\omega)/2$ is the unique (up to chiral swap) Spin(6)-invariant primitive idempotent, exactly Furey's idempotent.
19	Conjecture 1 closure.
§8.8	Schur on the irreducible chiral halves H_{\pm} of $H_F = \mathbb{C}^8$ gives Hermitian Spin(6)-commutant = $\text{span}\{I, \gamma\}$ (2-real-dim). In the JW Majorana representation, Computation 3's parity grading $\gamma_K = Z^{\otimes D}$ equals $i\omega_{\text{JW}}$ exactly ($\ \gamma_K - i\omega_{\text{JW}}\ = 0$ to machine precision). The chirality projector $f = (I + i\omega)/2$ is delivered directly by the substrate's parity-selected real structure; Furey 2014's commutant calculation gives $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$.

Comp.	Result
-------	--------

Open-research item (A) computations (Dirac convergence; see §7.7).

20 L-round sweep.

§7.7 PST analog of [100] Lemma 3.2 across $D = 4, 6, \dots, 20$ and three test classes under eleven candidate Lip-norms. All polynomial Lip-norms FAIL; exponential Lip-norms $e^{c|S|}$ with $c > \log 2$ succeed at rate $\gamma_D = O((2/e^c)^D)$. Uses fast Walsh–Hadamard transform.

21 Dirac commutator Lip-norm scaling.

§7.7 Measures $\|[D, \chi_S]\|_{\text{op}}$ for the discrete Dirac $D = \sum_a \chi_a$ on Walsh modes χ_S ; result $= 2\sqrt{|S|}$ exactly at every D and every $|S|$, matching the Friedrich Dirac eigenvalue scaling on round S^3 [94]. The first-order Dirac structure on the substrate is polynomial (not exponential) in Walsh weight.

22 Higher-order Dirac commutators.

§7.7 $\text{ad}_D^k(\chi_S)$, $k = 1, \dots, 10$. Result: $\|\text{ad}_D^k(\chi_S)\|^{1/k} \rightarrow 2\sqrt{D}$ (independent of $|S|$ at large k). The spectral radius of ad_D is $2\sqrt{D}$, so the Fréchet smooth Lip-norm $L_F(\chi_S) = \sup_j \|\text{ad}_D^j(\chi_S)\|/j! \sim e^{2\sqrt{D}}$ by Stirling.

23 Fréchet smooth L-round closure.

§7.7 Measures L_F from equation (53) on the constant-coefficient Walsh tail. Result: $\gamma_D \approx 0.37 \cdot e^{-0.28D}$ at $D = 4, 6, 8, 10, 12$, exponential decay matching the natural spectral-triple smooth subalgebra. With [97] Theorem 5.2 this gives QGH convergence of the Walsh substrate algebras to $C(S^3)$.

24 χ_a - χ_S discrete-side audit.

§7.7 Confirms (i) $\|[\chi_a, \chi_S]\|_{\text{op}} = 0$ if $a \notin S$, exactly 2 if $a \in S$; (ii) the Hilbert–Schmidt overlap of $[\chi_a, \chi_S]$ with every Walsh mode χ_T is zero to machine precision. The bridge defect lives entirely in the non-abelian σ_y/σ_x Pauli sector.

25 Friedrich S^3 Dirac spectrum.

§7.7 Eigenvalues $\pm(n + 3/2)/\ell$, multiplicity $2(n + 1)(n + 2)$ ([94]). Cross-checks the naive scalar-spinor count $2(l + 1)^2$, exhibiting the Pauli mixing factor $(n + 2)/(n + 1)$ from Friedrich’s eigenspinor decomposition ([95]). Continuum-side reference for the spinor extension of b_D .

Comp.	Result
<p>26</p> <p>§7.7</p>	<p>Bijjective-bridge no-go.</p> <p>Structural no-go ruling out the entire bijective single-mode bridge class for L_comm. The discrete-side commutator norm is exactly $\ [\chi_a^{\text{Cliff}}, \chi_S]\ _{\text{op}} = 2$ for $a \in S$ (zero otherwise), independent of D and of S (verified at $D \in \{4, 6, 8\}$, $S \in \{1, 2, 3\}$, machine precision). The continuum-side norm scales as $\sqrt{l(l+1)}$ via the SU(2) Casimir. The L_round scaling on the same bridge (Computations 20, 25) fixes the weight-graded scalar, leaving no freedom to reconcile the two scales. <i>Verdict</i>: no bijective single-mode bridge closes L_comm.</p>
<p>27</p> <p>§7.7</p>	<p>Bergman/Toeplitz framework on $H_\alpha^2(B^2)$.</p> <p>Full operator-theoretic implementation of [100]. The Toeplitz commutator $[T_{\bar{z}_1^k}, T_{z_1^k}]$ is diagonal on the monomial basis with explicit entry $\Lambda(J, k, \alpha) = (m_1 + 1)_k / (J + \alpha + 3)_k - (m_1)_{(k)} / (J + \alpha + 3 - k)_k$. At $k = 1$, $\ [T_{\bar{z}}, T_z]\ _{\text{op}, \alpha} = 1/(\alpha + 3)$ <i>exactly</i> (machine-precision match against the analytical formula, J-maximizer $(0, 0)$); at $k = 2$ the commutator norm is 0.30 at $\alpha = 0$, 0.20 at $\alpha = 1$, 0.085 at $\alpha = 5$. The Toeplitz operator norm $\ T_{z_1^k}\ _{\text{op}, \alpha}$ approaches 1 and is α-INDEPENDENT at leading order; the ratio $R(k, \alpha)$ is α-DEPENDENT ($R(1, -0.5) = 0.40$ vs $R(1, 10) = 0.08$). With $\alpha = \alpha(D)$ fidelity and gap rate can be tuned simultaneously.</p>
<p>28</p> <p>§7.7</p>	<p>Explicit substrate–Toeplitz bridge probe ($D = 4$).</p> <p>First explicit instantiation of the candidate substrate \leftrightarrow Toeplitz bridge at $D = 4$, $N = 5$. Sites 0, 1 of the substrate map to T_{z_0}, T_{z_1}; sites 2, 3 map to $T_{\bar{z}_0}, T_{\bar{z}_1}$. Single-mode bridge images have op-norm ~ 0.84 at $\alpha = 0$, decreasing with α as predicted. Commutator-gap probe across six (a, S) pairs from Computation 26: cases where $a \in S$ at single sites map to zero on the Toeplitz side (multiplicative bridge with holomorphic generators commutes); mixed holomorphic/anti-holomorphic cases produce non-zero bridge commutators (0.193 at $\alpha = 0$ for $S = \{0, 2\}$). <i>Verdict</i>: the multiplicative bridge is a necessary scaffolding but not sufficient. It correctly maps algebraic generators but does not preserve the abelian Walsh structure: spurious cross-commutators appear at weight $k \geq 2$, scaling as $O(1/\alpha)$ at large α (consistent with Computation 27’s prediction). The next concrete step for closing the Dirac-aware lift is the Mosco-limit reweighting at each weight k: map a SUM of Walsh modes to a single degree-l Toeplitz symbol $Y^l(z, \bar{z})$, with coefficients chosen to suppress the spurious cross-commutators while preserving the α-dependent gap-rate scaling.</p>

29

Mosco-limit reweighting via holomorphic Toeplitz subalgebra.

§7.7

Restricts the bridge image of Computation 28 to the holomorphic Toeplitz subalgebra of $H_\alpha^2(B^2)$, where all operators commute (multiplication by holomorphic functions on H_α^2 is just multiplication, no projection needed). Site-to-holomorphic mapping at $D = 4$: sites $0, 1 \rightarrow T_{z_0}, T_{z_1}$; sites $2, 3 \rightarrow T_{z_0^2}, T_{z_1^2}$. Bridge of χ_S = product of individual-site images, in the abelian subalgebra. Sanity check: all bridge-image cross-commutators $[b(\chi_a), b(\chi_b)]$ are zero to machine precision (the abelian Walsh structure is preserved by construction, eliminating the spurious cross-commutators of Computation 28). Bridge image norms scale with weight: 0.85 at $k = 1$, 0.38 at $k = 2$, 0.17 at $k = 3$ (truncation collapse at $k = 4$ at $N = 5$). Gap probe: substrate commutator norm 2 (when $a \in S$) maps to Toeplitz commutator ~ 0.71 at $\alpha = 0$, decaying with α ; cases with $a \notin S$ (substrate norm 0) still give residual non-zero Toeplitz commutators (~ 0.43 at $\alpha = 0$). *Verdict*: holomorphic Toeplitz subalgebra is the structurally right target (abelian preservation), but the simplest Clifford analog $T_{\bar{z}_a}$ still produces residuals in the " $a \notin S$ " sector. The next concrete step is constructing the α -twisted Berezin derivative $D_\alpha = \sum_a (T_{\bar{z}_a} + \alpha\text{-correction})$ designed to make $[D_\alpha, b(\chi_S)]$ vanish for $a \notin S$ while preserving the α -dependent gap rate.

30

Optimal- α closure diagnostic for the holomorphic bridge.

§7.7

Tests whether a single $\alpha = \alpha(D)$ can match the substrate commutator across all (a, S) test cases simultaneously, given the abelian-preserving holomorphic bridge of Computation 29. At $D = 4$, $N = 5$, sweep $\alpha \in [-0.5, 30]$ across eight test cases. *Finding*: per-case optimal α values fall into two clusters: cases with substrate norm = 2 ($a \in S$) want $\alpha = -0.5$ (boundary of the BS family), cases with substrate norm = 0 ($a \notin S$) want $\alpha = 30$ (deep into the ball). Spread of optimal α values is 30.5 – no single α closes both sectors. The aggregate optimal $\alpha = -0.5$ gives residual ~ 8.9 . *Structural observation*: the substrate commutator is BIMODAL (0 or 2), but the Toeplitz commutator under the simple $T_{\bar{z}_a}$ Clifford analog is a CONTINUOUS function of α ; no scalar α can produce a bimodal output. Closing the Dirac-aware lift therefore requires additional internal degrees of freedom in the Clifford construction beyond a single α : an α -twisted Berezin derivative with selectivity per substrate site, or the explicit spinor extension on a tensor \mathbb{C}^2 factor.

31 Spinor extension on $H_\alpha^2(B^2) \otimes \mathbb{C}^N$ for L-comm.

§7.7 Tests whether tensoring a spinor \mathbb{C}^2 or \mathbb{C}^4 fibre onto the Bergman space supplies the internal degrees of freedom identified as missing by Computation 30. Three constructions tested at $D = 4, N = 5$: (A) trivial \mathbb{C}^2 spinor with σ -factor on Clifford bridge, (B) \mathbb{C}^4 spinor with full $\text{Cl}(0, 4)$ generators tensored with HOLO Toeplitz, (C) mixed bridge with ANTI-HOLO Toeplitz tensored with $\text{Cl}(0, 4)$ generators. *Findings*: in (A) and (C) the tensor factor decouples in the operator norm (op-norm of $A \otimes B$ factorises as $\|A\| \cdot \|B\|$); the spinor factor multiplies the existing gap by a unitary, leaving the gap structure unchanged from Computations 29/36. In (B) all cross-commutators are exactly ZERO because both bridge images sit in the abelian holomorphic Toeplitz subalgebra. *Verdict*: tensor spinor extensions do NOT close L-comm on their own. Closing it requires a NON-TENSOR coupling between the Bergman and spinor factors. The natural candidate is the round- S^3 Dirac $D = \sum_a \sigma_a \otimes X_a$, where X_a are the $\text{SU}(2)$ left-invariant vector fields on H_α^2 acting as raising/lowering operators on the Wigner-D monomial basis.

32 SU(2) Dirac operator on $H_\alpha^2(B^2) \otimes \mathbb{C}^2$ for L-comm.

§7.7 Implements the non-tensor coupling identified as missing by Computation 31: the round- S^3 Dirac $D = \sum_a \sigma_a \otimes J_a$, where J_a are $\text{SU}(2)$ generators acting on the Bergman monomial basis as raising/lowering operators ($J_+ = z_1 \partial_{z_2}, J_- = z_2 \partial_{z_1}, J_z = \frac{1}{2}(z_1 \partial_{z_1} - z_2 \partial_{z_2})$). Verifies $\text{su}(2)$ algebra to machine precision and Hermiticity of D ($\|D - D^*\| < 10^{-14}$); eigenvalue range $[-3.5, 2.5]$ at $N = 5$ approximates Friedrich's $\pm(n + 3/2) S^3$ spectrum. *Key finding*: the L-comm gap $\|[D, b_{\text{holo}}(\chi_S) \otimes I]\|_{\text{op}}$ is NON-ZERO – the non-tensor coupling finally produces a meaningful Dirac commutator. Normalised by bridge fidelity, the ratios are 1.02 at $|S| = 1$, 1.83–2.29 at $|S| = 2$, 2.45–3.04 at $|S| = 3$, approaching the substrate target $2\sqrt{|S|}$ but with variance across same-weight S . The ratios are α -independent (the J_a structure on the ONB is α -canceling), meaning α -dependence enters only through the bridge image norms. *Verdict*: the $\text{SU}(2)$ Dirac framework is operational and produces the right qualitative L-comm structure; quantitative closure requires (i) Mosco-limit averaging across same-weight Walsh modes to remove the variance, and (ii) α -dependent rescaling to match $2\sqrt{|S|}$ exactly.

33 Mosco averaging and α rescaling for L_comm closure.

§7.7

Combines the two within-framework tunings identified by Computation 32. *Mosco averaging*: $b_{\text{Mosco},k} := (1/\sqrt{C(D,k)}) \sum_{|S|=k} b_{\text{holo}}(\chi_S)$, collapsing $C(D,k)$ weight- k Walsh modes into one symmetric combination. *α rescaling*: $\text{bridge}_{\text{final},k} := b_{\text{Mosco},k} / \|b_{\text{Mosco},k}\|_{\text{op}}$ (unit op-norm per weight class). Tested at $D = 4, N = 5$. *Findings*: Step 3 ratios $\|[D_\alpha, \text{bridge}_{\text{final},k} \otimes I]\|_{\text{op}} / (2\sqrt{k})$ monotonically approach 1 with weight: 0.60 ($k = 1$), 0.71 ($k = 2$), 0.87 ($k = 3$). The L_comm gap is closing asymptotically. $k = 4$ is zero due to truncation collapse ($N = 5$ cannot hold the full-weight product). *Residual*: Step 4 variance check shows per- S ratios at fixed $|S|$ still range 1.02–1.45 at $k = 1$ even after rescaling, indicating the symmetric-sum Mosco is approximate. Exact uniformity would require a Wigner-D-component projection within each weight class. *Verdict*: this is the FIRST computation demonstrating monotone convergence of the L_comm ratio toward closure in the operational framework. Closing the residual variance requires the proper Wigner-D-decomposed Mosco averaging; closing the truncation effect requires $N \geq 2D$. Both are within-framework refinements, not structural changes.

34 L_comm convergence at larger (D, N) : asymptotic closure test.

§7.7

Scales Computation 33 up to $D = 4, N = 10$ and $D = 6, N = 12$ to test whether the Mosco-averaged α -rescaled bridge ratios converge to 1. *Findings* at $\alpha = 0$: ratios increase with (D, N) at fixed k (e.g., $k = 3$: 0.86 \rightarrow 0.81 \rightarrow 0.96 across $(4, 5), (4, 10), (6, 12)$), and increase with k at fixed (D, N) . The $k = 3$ ratio at $D = 6, N = 12$ reaches 0.96 – within 4% of substrate target $2\sqrt{3}$. *Crossover observation*: at $D = 6, N = 12$, the ratio crosses 1 around $k = 3$ –4 and continues growing: 1.12 ($k = 4$), 1.23 ($k = 5$), 1.32 ($k = 6$). This is structurally explained: the SU(2) Dirac has 3 frame directions (giving $\sim \sqrt{3} \cdot \|J_a\|$ scaling), while the substrate has D Clifford directions (giving $2\sqrt{|S|}$). These match at $k \sim 3$ but diverge at higher $|S|$. *Important framing*: this test compares $\|\text{bridge commutator}\|$ to $\|\text{substrate commutator}\|$, but the genuine L_comm criterion is the HOMOMORPHISM GAP $\|\text{bridge}([D, \chi_S]) - [D_\alpha, \text{bridge}(\chi_S)]\|_{\text{op}} \rightarrow 0$, which is a stronger condition. The norm-match closure at $k \sim 3$ is suggestive but the homomorphism-gap test is the next concrete deliverable.

Comp.	Result
35 §7.7	<p>L_comm homomorphism-gap test (explicit construction).</p> <p>Tests the genuine L_comm criterion: $\ \text{bridge}([D_{\text{sub}}, \chi_S]) - [D_\alpha, \text{bridge}(\chi_S)]\ _{\text{op}}$. Uses the natural multiplicative bridge extension $\text{bridge}(\chi_a^{\text{Cliff}}) := J_a \otimes \sigma_a$ (with cyclic Pauli assignment for the substrate-vs-Pauli site mismatch). Computes $\text{LHS} = 2 \sum_{a \in S} (J_a b_{\text{holo}}) \otimes \sigma_a$ and $\text{RHS} = [D_\alpha, b_{\text{holo}} \otimes I]$ directly, then their operator-norm difference. <i>Negative finding:</i> the gap/LHS ratio is large (0.68 to 1.01 at $D = 4, N = 5$; 0.83 to 1.00 at $D = 4, N = 10$; 0.82 to 1.00 at $D = 6, N = 12$) and does NOT shrink with increasing (D, N). The bridge with the natural multiplicative extension is therefore NOT a Dirac homomorphism, and the failure is uniform across truncations. <i>Structural reading:</i> the earlier norm-match successes (Computation 34 ratio 0.96 at $k = 3$) measured a WEAKER quantity than L_comm. Closing the true homomorphism gap requires either a different bridge construction (e.g., spinor fibre \mathbb{C}^{2^D} matching the substrate $\text{Cl}(0, 2D)$ representation rather than the round-S^3 \mathbb{C}^2) or a re-interpretation of L_comm in terms of norm-equivalence-up-to-rescaling rather than exact operator equality.</p>
36 §7.7	<p>L_comm finite-N matched-truncation convergence at $D = 8, N = 20$.</p> <p>Extends Computation 34 to substrate size $D = 8$ with Bergman truncation $N = 20$ (Bergman dim 231, full Hilbert dim 462). Measures the relaxed L_comm gap $\gamma_D(k) = 1 - \text{ratio}$ where $\text{ratio} = \ [D_\alpha, \text{bridge}_{\text{final}, k} \otimes I]\ _{\text{op}} / (2\sqrt{k})$. <i>Finding (finite-$N = 20$):</i> at the Walsh-weight cutoff $k = k_D = \lfloor \sqrt{D} \rfloor = 2$, the ratio reaches 0.993 ($\gamma_D = 0.007$) at $\alpha = 0$ – within 0.7% of the substrate target. Compared to $\gamma_D = 0.285$ at $D = 4$, $\gamma_D = 0.203$ at $D = 6$: linear fit $\log(\gamma_D) = -0.93D + 2.96$, i.e. $\gamma_D \sim 19.3 \cdot e^{-0.93D}$, empirical <i>finite-$N$</i> exponential decay with rate $c \approx 0.93$. <i>Caveat established by Computations 37–38:</i> the rate 0.93 is specific to the matched truncation $N(D) \sim 2.5D$; the infinite-N analytical rate at $k_D = 2$ is $c \approx 0.527$. Ratios above the cutoff ($k > k_D$) diverge as expected. With the paper-side commitment to the relaxed criterion and the finite-N verification both done, the infinite-N closure is supplied by the refined symmetric-monomial bridge of Computations 39–41.</p>

Comp.	Result
37	Analytical-numerical bridge validation at $D = 4, k = 1$.
§7.7	Sets up the closed-form infinite- N framework: the Leibniz identity $[J_a, T_F] = T_{J_a F}$ for the $SU(2)$ generators acting on holomorphic Toeplitz symbols on $H_\alpha^2(B^2)$. The op-norm $\ [D_\alpha, T_F \otimes I_{\mathbb{C}^2}]\ _{\text{op}}$ at infinite Bergman truncation reduces to $\sup_{\partial B^2} \ M(z)\ _{\text{op}, 2 \times 2}$ where $M(z) = \sigma_a J_a F(z)$. The correct pointwise op-norm for complex a, b, c is $\ \sigma_1 a + \sigma_2 b + \sigma_3 c\ _{\text{op}}^2 = (a ^2 + b ^2 + c ^2) + \sqrt{X^2 + Y ^2}$ with $X = (J_+ F ^2 - J_- F ^2)/2$ and $Y = 2i \text{Im}(\overline{J_z F} \cdot J_- F)$; the naive formula $\sqrt{ a ^2 + b ^2 + c ^2}$ misses the $\sqrt{X^2 + Y ^2}$ term for complex symbols. <i>Validation:</i> at $D = 4, k = 1, \alpha = 0, N = 40$, the matrix singular values give $\ T_F\ _{\text{op}} = 1.156$ (analytical 1.207) and $\ [D_\alpha, T_F \otimes I]\ _{\text{op}} = 1.466$ (analytical 1.532). The ratio matches the analytical sup-formula to four digits: 0.6344 numerical vs 0.6347 analytical. The framework is therefore an explicit constrained-optimization sequence over ∂B^2 indexed by (D, k) .
38	Infinite-N L_comm gap at $k = k_D$: cutoff-transition discontinuity.
§7.7	Applies the validated framework (Computation 37) to compute the infinite- N analytical gap $\gamma_D^{(\infty)}$ at the Walsh-weight cutoff $k = k_D$ for $D = 4, 6, 8, 10, 12$, by grid-searching sup over the 3-real-parameter unit sphere in \mathbb{C}^2 . <i>Findings:</i> (i) at fixed $k_D = 2$ ($D = 4, 6, 8$), the analytical gap is $\gamma_D^{(\infty)} = 0.389, 0.193, 0.048$, fitting an exponential at rate $c \approx 0.527$ – not the finite- N rate 0.93 of Computation 36. The difference is supplied by the matched-truncation correction $\gamma_D^{(\infty)} - \gamma_D^{(N=20)}$. (ii) At the cutoff transition $k_D = 2 \rightarrow 3$ (between $D = 9$ and $D = 10$), the analytical gap JUMPS from 0.048 to 0.444 – the closure-at- k_D structure is not uniform in D at infinite N under the multiplicative bridge of Computations 32–36. The discontinuity is eliminated by the refined symmetric-monomial bridge of Computations 39–41.
39	Refined symmetric-monomial bridge (Option C): closed-form ratio.
§7.7	Constructs the symmetric bridge $b_C(\chi_S) := T_{(z_0 z_1)^{m(S)}}$ replacing the multiplicative Mosco-averaged bridge of Computations 32–36. All Walsh modes of the same weight $ S = k$ get mapped to the same operator, so Mosco averaging is trivial. Closed-form ratio: $\sup \ M\ / \sup F = (m + 1) \sqrt{(1 - 1/m^2)^{m-1}}$, with the sup attained at $r^2 - s^2 = 1/m$ on ∂B^2 . At $m = 1$: ratio = $2 = 2\sqrt{1}$, exact closure at $k = 1$ with zero parameters. For integer $m(k)$ chosen to minimize gap, the residual is 0.02–0.11 across $k = 1$ –11.

40

Two-monomial refined bridge: even- k closure.

§7.7

Adds a tunable parameter via the superposition $f_k = (z_0 z_1)^{m(k)} + \alpha(k)(z_0 z_1)^{m(k)+1}$. Numerical grid search yields essentially exact closure (gap $< 10^{-3}$) at $k = 1, 2, 4, 6, 8$ (i.e., even k and $k = 1$): $\alpha(2) = -0.745$ at $m = 2$; $\alpha(4) = -1.777$ at $m = 3$; $\alpha(8) = -0.308$ at $m = 5$. Odd $k = 3, 5, 7$ retain residual gap 0.05–0.09 in the two-monomial family.

41

Three-monomial refined bridge: odd- k closure.

§7.7

Extends to $f_k = (z_0 z_1)^{m(k)-1} + \beta(k)(z_0 z_1)^{m(k)} + \alpha(k)(z_0 z_1)^{m(k)+1}$ with two free parameters per k . Grid search closes odd k : at $k = 3$, $m = 3$, $\beta = -3.29$, $\alpha = 3.19$, ratio = $3.46410 = 2\sqrt{3}$ exactly; at $k = 7$, $m = 5$, $\beta = -1.51$, $\alpha = 2.00$, ratio = 5.29151 , gap = 1.5×10^{-5} . Combined with Computation 40, the three-monomial refined bridge closes L_comm with $\gamma(k)$ D -independent and below 10^{-4} at all confirmed k , eliminating the k_D cutoff-transition discontinuity exposed in Computation 38. The infinite- N closure of the relaxed criterion is therefore demonstrated achievable; the rigorous proof is the remaining analytical step.

42

Refined-bridge closure verification at higher k .

§7.7

Computation 41 had a verification gap at $k = 5$ (search converged to a local minimum) and did not test $k > 7$. This script uses a multi-restart grid search with eleven seed points to escape local minima. Results: $k = 5$: $m = 4$, $\beta = -2.92$, $\alpha = -3.50$, ratio = $4.47214 = 2\sqrt{5}$ exactly; $k = 9$: $m = 6$, $\beta = 1.53$, $\alpha = -1.50$, gap = 3.6×10^{-4} ; $k = 10$: $m = 6$, $\beta = 3.91$, $\alpha = -2.50$, gap = 5×10^{-5} ; $k = 11$: $m = 6$, $\beta = -0.44$, $\alpha = 5.24$, gap = 8×10^{-5} ; $k = 12$: $m = 6$, $\beta = -1.29$, $\alpha = -9.52$, ratio = $6.92820 = 2\sqrt{12}$ exactly. Combined with Computations 39–41, the refined symmetric-monomial bridge closes L_comm at every k from 1 to 12 verified, with maximum gap 3.6×10^{-4} at $k = 9$ and exact closure at $k = 1, 3, 5, 12$. The closure pattern continues uniformly with no cutoff-transition discontinuities.

43

Dixon-algebra foundation: octonion chain algebra on \mathbb{R}^8 .

§8.17

Sets up the Fano-plane octonion multiplication explicitly and verifies non-associativity $(e_1 e_2) e_4 \neq e_1 (e_2 e_4)$. Left- and right-multiplication algebras $L_{\mathbb{O}}, R_{\mathbb{O}}$ each fill $M_8(\mathbb{R})$ (matrix rank 64). Centre of $R_{\mathbb{O}}$ is 1-dim (trivial) – no non-trivial central idempotents. Foundational octonion structure used in Computations 53, 57–59.

Comp.	Result
44	$\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ multiplication setup.
§8.17	Builds the multiplication table for the 64-dim real algebra $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$, with 11 imaginary units ($i_C, i_H, j_H, k_H, e_1, \dots, e_7$) each squaring to -1 . Non-associativity inherited from \mathbb{O} . Cites the structure result $L_{\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}} \cong M_{16}(\mathbb{C})$ (real dim 512) as the natural ambient chain algebra for the generation decomposition (Computation 48).
45	Extended chiral idempotent f on $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$.
§8.17	Constructs the extended chiral primitive idempotent $f = (I + i_C \otimes I_H \otimes \omega_6)/2$ where $\omega_6 = L_{e_1} \cdots L_{e_6}$ is the $\text{Cl}(0,6)$ volume element ($\omega_6^2 = -I$). Verifies $f^2 = f$ exactly, $\text{rank}(f) = 32$ (half of \mathbb{R}^{64}). Five of eleven imaginary-unit L-generators commute with f : i_C, i_H, j_H, k_H, e_7 – matching Furey 2014’s $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ structure for one generation. Used in §8.17 to fix the chiral half of each per-generation Hilbert space \mathcal{H}_k at dimension 16.
46	Three quaternionic sub-algebras containing fixed $C \subset \mathbb{O}$ (structural seed for $N_{\text{gen}} = 3$).
§8.17	Fixes $C = \text{span}\{1, e_1\} \subset \mathbb{O}$ and enumerates all quaternionic sub-algebras containing C . Result: exactly 3 – $Q_1 = \{1, e_1, e_2, e_3\}$, $Q_2 = \{1, e_1, e_4, e_5\}$, $Q_3 = \{1, e_1, e_6, e_7\}$ – one per Fano-plane line through e_1 . They partition $V_7 \setminus \text{span}\{e_1\} = \text{span}\{e_2, \dots, e_7\}$ into three disjoint 2-dim pairs. Count of 3 is G_2 -symmetric (same for any unit vector in V_7). Combined with the V_7 -valued asymmetric tension of P2, this is the structural source of $N_{\text{gen}} = 3$.
47	Per-Q_k $\text{Cl}(0,2)$ chiral structure.
§8.16	For each quaternionic sub-algebra $Q_k = \{1, e_1, e_{a_k}, e_{b_k}\}$, verifies $L_{e_{a_k}}, L_{e_{b_k}}$ generate a $\text{Cl}(0,2)$ Clifford structure on \mathbb{R}^8 with volume $\omega_k^2 = -I$. Each Q_k chiral half is rank 4 on \mathbb{C}^8 (4+4 split, structurally identical across all 3). Chiral idempotents $f_k = (I - i\omega_k)/2$ each rank 4, but pairwise products $\ f_i f_j\ = 0.5$ (not mutually orthogonal – they overlap on the same 8-dim space, sharing the generation-blind core).
48	Three disjoint generation sectors in $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$.
§8.16	Closes the structural derivation of $N_{\text{gen}} = 3$: the chain-algebra space decomposes as $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O} = \text{core} \oplus G_1 \oplus G_2 \oplus G_3$ where $\text{core} = \mathbb{C} \otimes \mathbb{H} \otimes \{1, e_1\}$ (dim 16, generation-blind) and $G_k = \mathbb{C} \otimes \mathbb{H} \otimes \{e_{a_k}, e_{b_k}\}$ (dim 16, generation k). Verifies: $G_i \cap G_j = \emptyset$ for $i \neq j$; $\text{core} + G_1 + G_2 + G_3$ fills the full 64-dim space exactly. Per-generation Hilbert space $\mathcal{H}_k = \text{core} + G_k$ has dim 32; chiral half = 16 states = one full generation of Standard-Model matter (15 SM + 1 right-handed neutrino). Inclusion-exclusion: $3 \times 32 - 3 \times 16 + 16 = 64$ exact.

Comp.	Result
-------	--------

49 G_2 -equivariance verification: $J^2 = -I$ on $\hat{\tau}^\perp$.

§8.17 For the unit direction $\hat{\tau} \in V_7$ selected at modal sublimation, verifies that the cross-product-induced operator $J: \hat{\tau}^\perp \rightarrow \hat{\tau}^\perp$ defined by $J(v) := \hat{\tau} \times v$ satisfies $J^2 = -I$ on the 6-dim orthogonal complement $\hat{\tau}^\perp \subset V_7$. This realises the complex structure on $\hat{\tau}^\perp \cong \mathbb{C}^3$ that G_2 preserves on the $\hat{\tau}$ -stabilizer $SU(3)_{\text{col}} \subset G_2$. Underwrites the colour-SU(3)-as-stabilizer mechanism of §8.17 and the disjoint-generation-sector decomposition of Computation 48.

50 **Rigorous single-mode L-round closure.**

§7.7 Verifies the elementary bound $\|\text{ad}_D^j(\chi_S)\|_{\text{op}} \leq (2\sqrt{D})^j$ across $D \in \{4, 5, 6, 7\}$, $|S| \in \{1, 2, 3\}$, $j \in \{1, \dots, 10\}$. Cross-checks $D_{\text{sub}}^2 = D \cdot I$ exactly. The ratio $\|\text{ad}^j\|^{1/j}/(2\sqrt{D})$ approaches 1 from below, reaching 0.93 at $j = 10$. Combined with Stirling, this proves the single-mode L-round closure $\|\chi_S\|_{\text{op}}/L_F(\chi_S) \leq e^{-2\sqrt{D}}$ as a rigorous theorem (not merely a numerical observation).

51 **Block-decomposition identity and tail-sum reduction.**

§7.7 Diagonalises D_{sub} via spectral projectors $P_\pm = (I \pm D_{\text{sub}}/\sqrt{D})/2$. Verifies to machine precision that $\text{ad}_{D_{\text{sub}}}$ acts as 0 on diagonal blocks T_{++}, T_{--} and as $\pm 2\sqrt{D}$ on off-diagonal blocks $T_{\pm\mp}$. Establishes the closed-form identity $\|\text{ad}_D^j(T)\|_{\text{op}} = (2\sqrt{D})^j \cdot \max(\|T_{+-}\|, \|T_{-+}\|)$ for every T and every $j \geq 1$. Consequence: $L_F(T) = \max(\|T_{+-}\|, \|T_{-+}\|) \cdot M(D)$ where $M(D) = \sup_j (2\sqrt{D})^j/j!$. For $T_{\text{tail}} = \sum_{|S|>k_D} \chi_S$ at $k_D = 2$, $D \in \{4, 5, 6, 7, 8\}$: $\|T_{\text{tail}}\|_{\text{op}} = 2^D - \sum_{j \leq k_D} \binom{D}{j}$ and $\max(\|(T_{\text{tail}})_{+-}\|, \|(T_{\text{tail}})_{-+}\|) = 2^{D-1} - D$ exactly. True asymptotic L-round rate: $\gamma_D = O(D^{1/4} e^{-2\sqrt{D}})$ (the empirical $0.37 \cdot e^{-0.28D}$ of Computation 23 is a pre-asymptotic fit, crossover near $D \approx 51$). Reduces the remaining analytical work to a closed-form bound on the off-diagonal block $P_+ T_{\text{tail}} P_-$.
