Precausal Substrate Theory (PST) proposes that spacetime, causality, matter and the fundamental forces are not the bedrock of reality; they arise from logic alone.

The sole primitive and first postulate is property differentiation (P1): the capacity for distinctions to exist at all. It carries an asymmetric vector tension (P2); and when that tension crosses a threshold, modal sublimation (P3) produces instantiated geometry. From these three postulates, together with the established methods of noncommutative geometry, PST builds a common foundation beneath general relativity and quantum mechanics. It recovers the Standard Model gauge structure and a three-generation count of matter. It predicts a Casimir-force correction of derived form, its size below present reach. And it derives a Higgs self-coupling consistent with the measured Standard-Model value at 1.4 sigma.

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.



ISBN 978-94-6567-147-5

9 789465 671475

Second Edition – EU version

Precausal Substrate Theory

Ralf Meelker

Second Edition

Precausal Substrate Theory

On the Nature of Reality, as I see it

Ralf Meelker

Second Edition

Precausal Substrate Theory

On the Nature of Reality, as I see it

Ralf Meelker

2026

Second Edition

Copyright © 2026 by Ralf Meelker
All rights reserved
Authored by Ralf Meelker
Printed by Pumbo.nl BV
Published by Nashua BV

Precausal Substrate Theory
First Edition, released June 2026
First Edition, second print July 2026
Second Edition, released October 2026

Companion site https://meelker.nl/pst-se
Computations https://meelker.nl/pst-se/computations

ISBN 978-94-6567-147-5 (EU paperback version)
ISBN 978-90-9043-536-7 (US paperback version)
NUR 925

Preface

In this book I pursue a single idea as far as it will go, and Precausal Substrate Theory is where it leads: that spacetime, causality, matter, and the fundamental forces are not the bedrock of reality; they arise from logic alone. The sole primitive and first postulate is property differentiation (P1), the bare capacity for distinctions to exist at all; it carries an asymmetric vector tension (P2), and when that tension crosses a threshold, modal sublimation (P3) produces instantiated geometry.

When I started, I did not set out to write a final theory, only to see whether such a precausal substrate could carry the weight physics usually places on a presupposed spacetime. It carried more than I expected. From these three postulates, together with the established methods of noncommutative geometry, PST builds a common foundation beneath general relativity and quantum mechanics, and recovers the Standard Model gauge structure and a three-generation count of matter. It predicts a Casimir-force correction of derived form, its size below present reach. It also yields a derived Higgs self-coupling that matches the measured value at 1.4 sigma, resting on one stated premise, that the Standard Model is PST's low-energy theory, which a renormalisation-group test in the book passes.

For transparency, the modelling choices that enter the quantitative results, and the dependence on noncommutative geometry, are set out in full in the book, with all computations on the companion website.

The subtitle says "as I see it", but the aim is not private: my theory is set down plainly enough that another reader can see it the same way. Most striking is how much follows from how little; I hope it encourages others to research further.

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.

Ralf Meelker, October 2026

Contents

Preface **iii**

Summary **xix**

1 Introduction **1**

1.1 The incompatibility of quantum mechanics and relativity . . . 1

1.2 Assumed backgrounds and the missing question 2

1.3 The precausal starting point 3

1.4 The case for ontological minimality 4

1.5 The foundational object: the KO-dimension-ten triple 5

1.5.1 The substrate–Higgs identification 9

1.6 Principal consequences 10

1.7 The foundational postulates of PST 13

1.7.1 Derived statement: the Laplace-type projection 16

1.8 Structure of this book 17

1.9 Summary 19

2 Property Differentiation: The Sole Primitive **21**

2.1 Differentiation by non-identity alone 21

2.2 The substrate triple (D, δ, v) 22

2.3 The direction algebra V_7 23

2.4 Configurations and configuration space 24

2.5 The canonical measure on configurations 24

2.5.1 Theorem (Canonical Measure) 26

2.6 Summary 27

3 Asymmetric Vector Tension: The Preferred Direction **29**

3.1 The structural origin of asymmetric tension 29

3.2 The tension functional and the direction algebra V_7 30

3.3 Epistemological status of the formalism 31

3.4 Complement asymmetry and the substrate's directional bias $\hat{\tau}$ 32

3.5 The modal and ontological status of $T(C)$ 33
3.6 Philosophical lineage: Leibniz, Spencer-Brown, Hegel 40
3.7 Summary 41

4 Modal Sublimation: How the Possible Becomes Real 43

4.1 Why a modal threshold must exist 43
4.2 The boundary condition on $|T(C)|$ 47
4.3 The order parameter and its ontological status 48
4.4 Variational formulation of τ 49
4.5 Stability analysis and the two regimes 57
4.5.1 Structural versus spontaneous asymmetry 59
4.6 The bifurcation point and the derivation of τ 61
4.6.1 Notation: which threshold τ names 63
4.7 Summary of the threshold half 63
4.8 The character of the threshold 65
4.9 Categorical formulation: the sublimation functor 68
4.10 Explicit construction of Φ 70
4.11 Uniqueness of Φ and the origin of diffeomorphism invariance . 75
4.12 Lorentzian signature from Goldstone hyperbolic well-posedness 77
4.13 Summary 81

5 Coinstantiation of Spacetime and Causality 83

5.1 Causality as coinstantiated structure 83
5.2 Mathematical formulation of coinstantiation 84
5.3 The direction of time as a structural consequence 85
5.4 The causal order and the degenerate vacuum 87
5.5 Category error and the dissolution of first-cause questions . . 87
5.6 Contrast with competing approaches 89
5.7 Summary 90

6 The Emergence Chain: Eight Steps to a Universe 93

6.1 Overview and the eight-step summary table 93

6.2 The character of the chain 93

6.3 Step 1 to Step 2: from differentiation to tension 96

6.4 Step 2 to Step 3: from tension to threshold 96

6.5 Step 3 to Step 4: from threshold to sublimation 97

6.6 Step 4 to Steps 5 & 6: sublimation to spacetime & causality . 98

6.7 Steps 6 and 7: from causality and geometry to matter 99

6.8 Step 8: from geometry and the vacuum to angular momentum 100

6.9 The chain as a closed logical structure 101

6.10 Summary 103

7 Energy and Matter as Projected Tension 105

7.1 Results imported from the technical core 105

7.2 The source problem in General Relativity 107

7.3 The stress-energy tensor as projected tension 107

7.4 The tensor projection $\Pi_{\mu\nu}$ 108

7.4.1 Why the obvious completion fails 109

7.4.2 The separation dyad and its trace 109

7.4.3 What fixes the equation of state 110

7.4.4 The weight is not free either 110

7.4.5 What is not derived 112

7.4.6 What the matter sector can and cannot be 115

7.5 The field equations as a conservation statement 116

7.5.1 Why both routes are kept 123

7.6 The F–S dichotomy: free energy versus action 125

7.7 The projected action and the Einstein–Hilbert term 126

7.8 Seeley–DeWitt argument for the projected action 128

7.9 Summary 135

8 Gravity from the Projected Geometry 137

8.1 Spacetime curvature 137

8.2 The weak-field limit and the recovery of Newtonian gravity . . 138

8.2.1 Post-Newtonian parameters and scalar-mode decoupling . . . 139

8.3 The strong-field limit: horizons, and why light cannot leave . . 140

8.3.1 A horizon is a threshold surface, not a trap 141

8.3.2 Why light cannot leave, when nothing is pulling on it 141

8.3.3 Where the projected field equations stop 142

8.3.4 What the null threshold fixes about light 143

8.3.5 The Planck spectrum, assembled 144

8.3.6 What the ladder already settles about the photoelectric effect 146

8.3.7 The Hawking temperature, and four things not claimed 148

8.4 Stress-energy as projected tension 150

8.5 Newton's constant as a consistency relation 151

8.6 Uniqueness of $G_{\mu\nu}$: Lovelock's theorem 152

8.7 The equivalence principle from single-source projection 154

8.8 The gravitational constant as a unit conversion 154

8.9 G from the projection scale: form and magnitude 155

8.10 Kernel normalisation: the finite-internal-factor trace 157

8.11 Summary 158

9 Vacuum Structure and Angular Momentum 159

9.1 The cosmological constant: substrate vs emergent energy . . . 159

9.1.1 What $\mathcal{F}\big|_{\mathcal{V}}$ is, and is not 159

9.1.2 Mechanism A: bare Goldstone zero-point on the slice $\mathcal{V}_{\dot\eta} \cong S^1$ 160

9.1.3 Mechanism B: the ongoing-sublimation source 161

9.2 The degenerate vacuum manifold and angular momentum . . 163

9.3 The incomplete vacuum of standard physics 166

9.4 Angular momentum as Noether charge 167

9.5 Angular momentum and the Goldstone–geodesic connection . 169

9.6 Summary . . . 173

10 Recovering Quantum Mechanics 175

10.1 Structural and specific predictions . . . 175
10.2 The Born rule from the Bernoulli measure . . . 176
10.2.1 Two non-contextualities, and which one is assumed . . . 178
10.3 The full measurement postulate from threshold-crossing . . . 179
10.3.1 The threshold-crossing event as collapse . . . 179
10.3.2 Probability of selection . . . 180
10.3.2.1 Closure (1): Explicit measurement-coupling Hamiltonian . . . 181
10.3.2.2 Closure (2): Decoherence inheritance . . . 182
10.3.2.3 Closure (3): threshold-crossing uniqueness . . . 183
10.3.2.4 Theorem (threshold-crossing uniqueness) . . . 183
10.3.2.5 Closure (4): agreement with the Born rule . . . 185
10.3.3 Why this is not the many-worlds picture . . . 188
10.3.4 Why this is not an objective-collapse theory . . . 189
10.4 The canonical experiments . . . 190
10.4.1 The double slit and which-path information . . . 191
10.4.2 Delayed choice and the quantum eraser . . . 194
10.4.3 Stern–Gerlach and the definite eigenvalue . . . 196
10.5 Composite systems, the reduced state and no-signalling . . . 199
10.5.1 The bipartition is of the property domain, not of space . . . 199
10.5.2 Theorem (tensor factorisation over a bipartition) . . . 200
10.5.3 The reduced state and local observables . . . 201
10.5.4 No-signalling across the bipartition . . . 202
10.5.5 Bell correlations and emergent spatiality . . . 203
10.6 Canonical commutation relations from the Boolean gradient . . . 205
10.7 Quantum fluctuations . . . 206
10.8 Summary . . . 208

11 Spin and Statistics from Boolean Exclusion **209**

11.1 The fermionic sector: structure and the open problem 209

11.2 Fermionic excitations: the CAR algebra identification 210

11.3 Spin-statistics from Boolean exclusion (structural theorem) . . 212

11.3.1 Structural theorem (Boolean spin-statistics) 212

11.3.2 Consequence: half-integer excitations are fermionic 213

11.3.3 Relation to the standard QFT spin-statistics theorem 213

11.3.4 Closure (1): CAR convergence 214

11.3.5 Theorem (CAR Mosco convergence) 214

11.3.6 Closure (2): Wightman axioms 216

11.3.7 Closure (3): the Dirac equation 218

11.4 Summary 220

12 Quantum Fields and the Higgs Quartic **221**

12.1 The projected field theory: emergence of quantum fields . . . 221

12.2 Interacting vacuum and M_*-regulated renormalisation 225

12.2.1 Scope of the cancellation argument 235

12.2.2 Candidate substrate-side identification of Z^2 240

12.3 Closure via the binary-gap-tempered partition function 241

12.3.1 An independent substrate-side derivation of e^{-1} 241

12.3.2 The substrate Higgs Hamiltonian 242

12.3.3 What the identification carries, and what it does not 243

12.3.3.1 Notation: the KO subscript 243

12.3.4 The bridge identification: seven-step proof sketch 244

12.3.4.1 Theorem sketch ($Z^2 \approx \mathrm{e}^{-1}$ bridge) 246

12.4 The spectral-action route to the Higgs quartic 248

12.4.1 Conjecture (spectral-action Higgs-quartic identity) 248

12.4.2 The Yukawa as a fourth root of the measure 250

12.4.3 First-cut inner-fluctuation calculation 250

12.5 Partition-function Chamseddine–Connes framework 252

12.5.1 The eleven statements of the framework (F1–F11) 252

12.5.2 The bridge premise and its reduction . . . 253
12.5.2.1 Reduction of (B): the substrate side closed . . . 254
12.5.3 What the four steps establish, and what remains . . . 256
12.5.4 Substrate-side closure; the open matching step . . . 258
12.6 The framework in full: F1–F11 . . . 258
12.6.1 Definitions (F1–F3) . . . 258
12.6.1.1 Definition F1 (PST partition-function spectral triple) . . . 258
12.6.1.2 Definition F2 (partition function of an observable) . . . 258
12.6.1.3 Definition F3 (Matched-scaling exponent) . . . 259
12.6.2 Propositions (F4–F7) . . . 259
12.6.2.1 Proposition F4 (PST EFT below M_*) . . . 259
12.6.2.2 Proposition F5 (Wilsonian matching at M_*) . . . 259
12.6.2.3 Proposition F6 (PST UV Higgs quartic) . . . 260
12.6.2.4 Proposition F7 (Substrate UV-IR scale separation) . . . 260
12.6.3 Theorems (F8–F11) . . . 260
12.6.3.1 Theorem F8 (substrate Higgs Hamiltonian) . . . 260
12.6.3.2 Theorem F9 (tempered Bernoulli partition function) . . . 261
12.6.3.3 Theorem F10 (unique soft mode; total reduction) . . . 261
12.6.3.4 Main theorem of the framework . . . 262
12.6.3.5 Theorem F11 ($Z^2 \to \mathrm{e}^{-1}$ in the substrate) . . . 262
12.6.4 Remark on what the framework achieves . . . 263
12.7 Status of the derivations . . . 263
12.8 Summary . . . 266

13 Gauge Structure, Quantisation, and the Standard Model 267

13.1 From U(1) to the full Standard Model gauge group . . . 267
13.2 Path-integral quantisation of the projected theory . . . 268
13.3 Local gauge symmetry: the photon as a connection . . . 269
13.3.1 Minimal coupling is forced, and this is the algebra . . . 270
13.4 $\mathrm{SU}(2)_L$ from threshold topology: the electroweak doublet . . . 271
13.4.1 Two descriptions of one would-be Goldstone phase . . . 272
13.4.2 Relation to existing electroweak approaches . . . 274

13.4.3 The doublet does not fix chirality . . . 274
13.5 Colour SU(3)$_c$: octonionic origin and Casimir decomposition . 274
13.5.1 Octonionic origin: the G_2 stabiliser of a V_7 direction . . . 275
13.5.2 Rigorous Casimir decomposition . . . 277
13.5.3 What the decomposition does and does not settle . . . 278
13.5.4 The strong sector: what is derived and what is inherited . . . 279
13.6 Consistency of the octonionic interpretation . . . 280
13.6.1 Embedding-independence of the structural results . . . 281
13.6.2 Unitary covariance of the order-one selection . . . 281
13.6.3 Structural scope under the octonionic embedding . . . 282
13.7 The internal algebra A_F derived from P1–P3 . . . 282
13.7.1 Piece 1: the chirality grading . . . 283
13.7.2 Piece 2: f as the canonical idempotent . . . 283
13.7.3 Piece 3: SU(3)$_c \times$ U(1) from the chain algebra . . . 284
13.7.4 Piece 4: SU(2) from the Lorentzian Dirac sector . . . 285
13.7.5 Piece 5: Dixon hyperspinor synthesis . . . 286
13.7.6 Conjecture 1, closed (within the Dixon synthesis) . . . 286
13.7.7 Cl$(0,6)$ representation theory . . . 288
13.7.8 The minimal left ideal . . . 289
13.7.9 Furey's chain-algebra action and the commutant . . . 289
13.7.10 Cl$(1,3) \cong M_2(\mathbb{H})$ and the internal SU(2) . . . 290
13.7.11 Dixon synthesis: $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ . . . 291
13.7.12 The algebra identity $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ . . . 292
13.7.13 Three generations from six SU(3) triplets . . . 292
13.7.14 The gauge group as substrate automorphisms . . . 293
13.8 Gauge anomaly cancellation as a structural check . . . 294
13.9 U(1)$_Y$ hypercharges compatible with PST structure . . . 296
13.10 The Standard Model Lagrangian as leading-order projection . 298
13.11 Summary . . . 299

14 Fermion Content: Generations, Yukawa and Chirality **301**

14.1 Fermionic representation content under Π 301

14.2 The three-generation count from the V_7 directional content . . 304

14.2.1 The structural chain 304

14.2.2 Three disjoint generation sectors 305

14.2.3 Colour and generation: two guises of one three 306

14.3 Yukawa and CKM: the structural-scope theorem 308

14.4 Chirality of $\mathrm{SU}(2)_L$ as a structural consequence 310

14.5 Summary 311

15 Grounding d_0 and the Casimir Prediction **313**

15.1 Modal sublimation and the substrate scales 313

15.2 The modal scale and the electroweak identification 314

15.2.1 Relation to other approaches to the Higgs mechanism 317

15.3 Grounding the scale d_0: the coherence length 317

15.3.1 Two distinct scales: grain versus coherence length 318

15.4 Vacuum phenomena and the Casimir effect 319

15.5 Casimir detectability at each separation regime 321

15.6 The parameter-free exponent and falsifiability 324

15.6.1 Kernel-shape sensitivity and the d_0 bound 325

15.6.2 The Gaussian as the substrate's kernel 326

15.7 Falsifiability and the experimental programme 328

15.8 Observational constraint on d_0 from Casimir data 330

15.9 Correction budget for the Casimir measurement 331

15.9.1 Predicted deviation across separations 332

15.10 Experimental status of M_* and the ADD comparison 332

15.11 Wilson coefficients: tree-level custodial ΔT protection 335

15.12 Summary 337

16 Cosmology: Acceleration from Ongoing Sublimation 339

16.1 PST's reframing of the standard cosmological picture 339

16.2 The Cosmological Principle from the Bernoulli measure 339

16.3 Horizon and flatness problems dissolved 341

16.4 Ongoing instantiation: why it does not dilute 341

16.5 Equation of state and the sign of the cosmological constant . . 342

16.6 Friedmann equations and cosmic acceleration 343

16.7 The rate problem 344

16.8 Observational commitments 345

16.9 What would refute the constancy, and what would not 346

16.9.1 The falsifier, stated sharply 347

16.9.2 Where the measurements currently stand 348

16.9.3 Why PST cannot retreat, and what that costs 351

16.10 Comparison with steady-state cosmology 352

16.11 Summary 353

17 Standing Among Existing Frameworks 355

17.1 PST as foundation, not competitor 355

17.2 Comparison with established physics frameworks 355

17.3 String theory: the scope contrast 360

17.4 Non-Commutative Geometry 361

17.5 Other approaches to the Higgs mechanism 362

17.6 Inflationary Cosmology and the Multiverse 364

17.7 Spencer-Brown and Distinction-Based Ontologies 366

17.8 Wheeler's "It from Bit" and the Mathematical Universe . . . 367

17.9 Structural Realism and Bohm's Implicate Order 368

17.10 How the claims in this book are meant to be checked 369

17.11 Summary 371

18 Convergence of the Projected Geometry **373**

18.1 Mosco convergence: theorem and assumptions 373
18.1.1 Setup: the substrate sequence and matched scaling 373
18.1.2 The seven assumptions A1–A7 374
18.1.3 The Mosco-convergence theorem 376
18.1.3.1 Theorem (Mosco convergence, under A1–A7) 376
18.1.3.2 Proof sketch of the Mosco theorem 377
18.1.4 Numerical verification on finite models 378
18.1.5 What the theorem establishes 378
18.1.6 Sharpening A7: a uniform-rate inequality 379
18.1.6.1 Lemma (A7 reduction) 379
18.1.7 The physical content of the bound 380
18.1.8 K-uniformity of the constant 381
18.1.9 The analytic content of the rate hypothesis 381
18.1.10 Numerical verification of the rate 384
18.2 The K-uniform rate: proof and three closures 385
18.2.1 Setup: the domain K and the empirical aggregate 385
18.2.2 Step 1: Taylor expansion 386
18.2.3 Step 2: Expectation under ξ 386
18.2.4 Step 3: Empirical average over the realisation map 386
18.2.5 Step 4: K-uniform bounds 387
18.2.6 Step 5: K-uniform rate 388
18.2.7 Caveats and what remains 388
18.2.8 Closure (a): Explicit $\Pi_n^\dagger$ construction 389
18.2.9 Closure (b): Thin-support d_0-corrections 390
18.2.10 Closure (c): Riemannian extension on $M = \mathbb{R} \times S^3$ 390
18.2.11 Explicit constants: Berry–Esseen and finite elements 391
18.3 Dirac-operator convergence: a single analytic estimate 392
18.3.1 The three ingredients of the Dirac estimate 392
18.3.2 Ingredient 3, what is exact at finite substrate size: the Clifford algebra, the multiplicities and the bridge map 393

18.3.3 Ingredient 3, what remains: the commutator residual as a Sobolev-tail estimate . . . 393
18.3.4 Step 5, the discrete-derivation defect: why no citation closes it, and what does . . . 394
18.3.4.1 Section round-trip in Gromov–Hausdorff distance . . . 395
18.3.5 Dirac-aware lift to substrate-to-continuum metric . . . 396
18.3.6 Why the relaxed criterion suffices for PST . . . 399
18.3.6.1 Caveat on the rate . . . 400
18.3.7 Refined bridge: the Lemma 5–8 closure of L_comm . . . 400
18.4 Dirac convergence: the L_round and L_comm lemma chain . . . 402
18.4.1 The L_round side: the Gromov–Hausdorff round-trip . . . 402
18.4.2 The L_comm side: the refined-bridge lemma chain . . . 405
18.4.2.1 Refined bridge, symmetric-monomial construction . . . 405
18.4.2.2 Refined bridge: the ratio match (Lemma 5) . . . 406
18.4.2.3 Refined bridge, Berezin compatibility (Lemma 6) . . . 407
18.4.2.4 Refined bridge: finite-N closure rate (Lemma 7) . . . 408
18.4.2.5 Refined bridge: spectral-action invariance (Lemma 8) . . . 410
18.5 Summary . . . 411

19 Resolved Problems and the Open Frontier 413

19.1 The resolved sectors . . . 413
19.1.1 The foundational object . . . 413
19.1.2 Emergent geometry and the Lorentzian signature . . . 414
19.1.3 Gravity . . . 414
19.1.4 Cosmology . . . 415
19.1.5 Quantum mechanics . . . 415
19.1.6 The Standard Model . . . 416
19.2 The Dirac-convergence step . . . 417
19.2.1 Reducing the bound to ladder commutators . . . 418
19.2.2 Asymptotic Toeplitz-symbol closed form for $c_1(k)$ . . . 423
19.2.3 Analytical uniform-in-(k, N) bound . . . 423
19.3 Octonionic-SU(3) structure inside $\mathrm{Cl}(0, 6)$ . . . 425

19.4 Matter, and the theory's empirical exposure 426
19.4.1 Boolean occupation makes spin-statistics structural 426
19.4.2 A 6-sphere vacuum quantises angular momentum 427
19.4.3 The three places the theory can be measured 427
19.4.4 Where PST matches general relativity, on purpose 429
19.4.5 What the theory does not predict: the fermion masses and mixings, and why that is a theorem 429
19.5 The A_F algebra identity 430
19.6 Robustness of the structural-scope theorem 431
19.7 The substrate-side identification of Z^2 431
19.7.1 The CC inner-fluctuation bridge is obstructed 432
19.7.2 An alternative substrate-to-SM bridge 432
19.7.3 The substrate side of the bridge, closed 433
19.7.4 The matching, by wave-function renormalisation 433
19.7.5 Empirical-precision test 433
19.7.6 The bridge proof sketch 434
19.8 Standing among substrate theories 435
19.9 The remaining open frontier 436
19.9.1 The structural front 437
19.9.2 Localising an apparatus in the property domain 438
19.9.3 The black-hole interior 441
19.9.4 Beyond the two quantitative bounds 443
19.9.5 Bridge Premise (B), reduced 446
19.9.6 The full-SM extension, settled in the negative 447
19.10 Summary 448
20 Conclusion **449**
20.1 The shape of the argument 449
20.2 The central claim, stated precisely 449
20.2.1 Standing among substrate theories 450
20.2.2 The Dirac-convergence constant 450

20.3 What the emergence chain establishes 451
20.4 The open questions PST closes 452
20.5 The unification of quantum mechanics and general relativity . 454
20.6 The quantitative standing of the theory 454
20.7 The scope and limits of the theory 457
20.7.1 What is foundational, and what is inherited 457
20.7.2 Structurally contingent versus derived 461
20.8 Open questions and the next research 462
20.9 On intelligibility and necessity 463
20.10 Summary 464

Appendix A: Notation Conventions 467

Appendix B: Computation Scripts 489

Appendix C: Claims Status 639

Appendix D: Comparison to Other Theories 669

Appendix E: Bibliographic References 683

Summary

Precausal Substrate Theory (PST) proposes that spacetime, causality, matter, and the fundamental forces are not the bedrock of reality; they arise from logic alone. The sole primitive and first postulate is **property differentiation (P1)**: the capacity for distinctions to exist at all. The second postulate is that property differentiation carries an **asymmetric vector tension (P2)**. The third postulate is **modal sublimation (P3)** past a Landau–Ginzburg threshold, producing instantiated geometry along the substrate's net direction; causality is born together with geometry. These three are the theory's only foundational postulates; the later structural results follow from them together with a *single* extremal selection principle (maximality over the Hurwitz classification, applied to the direction algebra and to the hyperspinor envelope), which is itself a consequence of P1 rather than a further commitment (§1.5, Computation 168), and with the standard Chamseddine–Connes noncommutative-geometry machinery (inner fluctuations, the spectral action) applied to the substrate spectral triple that P1–P3 deliver.

From these three postulates a concrete mathematical structure follows. PST delivers general relativity, quantum mechanics, the Standard Model gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$, fermionic matter with structural spin-statistics, and the three-generation count of matter, $N_{\mathrm{gen}} = 3$, with the Higgs identified as the projection of the substrate's modal order parameter. The fermion mass and mixing spectrum, the Yukawa hierarchy and the Cabibbo–Kobayashi–Maskawa and Pontecorvo–Maki–Nakagawa–Sakata mixing matrices, is by construction not among these structural outputs. That placement is itself a theorem, not a concession: the structural layer fixed by P1–P3 is generation-symmetric at leading order, so the observed hierarchy and mixing cannot arise from it and must enter through the contingent configuration content $T(C)$. PST predicts a Casimir correction with a parameter-free, signature-distinguishable d^{-6} exponent, its amplitude set by the coherence scale d_0. Current data bound $d_0 \lesssim 2.3$ nm; the theory's own canonical normalisation fixes it at 2.2 am (Computation 201), so the predicted deviation is exact in form and far below

present reach. A null result at any reachable precision is what the theory expects, and a nanometre-amplitude signal would refute the normalisation rather than confirm the book. Its other quantitative result, a derived value for the Higgs self-coupling, is consistent with the Standard Model at 1.4 sigma on current inputs (Computation 202). Among twenty-seven surveyed substrate theories (PST and twenty-six others), PST alone delivers a pre-geometric substrate, emergent spacetime, Standard-Model gauge ingredients, and the three-generation count of matter in one framework.

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 emergence chain

The three postulates compose into a single derivation chain rather than three independent assumptions. P1 supplies a Boolean configuration space $\mathcal{P}(D)$, the powerset of D substrate sites, with cardinality 2^D, on which a Bernoulli product measure μ realises the principle of equal a priori weighting of distinctions. P2 endows $\mathcal{P}(D)$ with an asymmetric vector tension T: each configuration carries a signed contribution $T(C)$ measuring its directional misalignment with the substrate's net polarity. P3 says that when the integrated tension over a coherent region exceeds the Landau–Ginzburg threshold τ, the substrate spontaneously breaks its Boolean symmetry and crystallises a six-sphere of minima, $\mathcal{V} \cong S^6 \subset V_7$, whose 2-plane slice is the familiar ring. This crystallisation event is *instantiated geometry*. The modal angle θ becomes the (eaten) Goldstone mode whose phase aligns into the Higgs vacuum, and the radial fluctuation becomes the physical Higgs scalar. The post-sublimation manifold $M = \mathbb{R} \times S^3$ inherits its Lorentzian signature and the direction of causality from the symmetry-breaking pattern, and its macroscopic dimension $n = 4$ from the minimal split of the spectral triple's total dimension (below). Causality is born together with geometry, not imposed on it. Everything that follows in this Summary operates within this post-sublimation manifold, with the substrate retained as the ultraviolet (UV) layer above the modal scale

$M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV.

The Core Mathematics

The substrate is a real Connes spectral triple of real K-theory (KO) dimension six, the mod-8 index that fixes which real structure a spectral triple can carry, with sign pattern $(+,+,-)$ (Computation 3), the internal value of the Connes Standard-Model construction, whose product with the spacetime factor's KO-4 lands the total on the Lorentzian-compatible two. Its Mosco limit, the limit in which the substrate's discrete energy functional becomes the continuum's and carries its operators with it, yields spacetime $M = \mathbb{R} \times S^3$ with KO-dimension four (Computation 4); the product $M \times F$ has total KO-dimension $10 \equiv 2 \pmod 8$, the Connes Standard-Model value, with the macroscopic dimension $n = 4$ fixed by the minimal Connes-spectral-triple split $(4 + 6 = 10 \equiv 2 \pmod 8)$ given the verified KO-6 substrate. The substrate's parity grading equals the chirality grading $i\omega$ of $\mathrm{Cl}(0, 6)$, the Clifford algebra on six anticommuting generators, exactly (Computation 19), delivering the grading on which Furey's construction rests [121]. Furey's chain-algebra construction on $\mathbb{C} \otimes \mathbb{O}$, whose left-multiplication algebra is $\mathrm{Cl}(0, 6)$, delivers $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation of fermion content under $\mathrm{SU}(3)_c$. The emergent 4-d Lorentzian Dirac sector has local Clifford algebra $\mathrm{Cl}(1, 3) \cong M_2(\mathbb{H})$ (Cartan classification [153, 154]), so right-multiplication by $\mathbb{H}$ on the spinor module is an internal symmetry (Dixon 2010 [146], Computation 66), generating $\mathrm{Sp}(1) \cong \mathrm{SU}(2)$. Combined via Dixon's hyperspinor framework $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$, these structures synthesise to the full Standard-Model internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, whose unimodular unitaries are $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$.

The Unified Physics

The same spectral-triple structure delivers the four otherwise-disjoint foundations of contemporary physics: gravity, the quantum, the gauge forces, and matter. *General relativity*: the Einstein–Hilbert action arises as the leading term of the Connes spectral action $\mathrm{Tr}\, f(D/\Lambda)$, with Newton's constant a consistency relation among (m_h, M_P, v). *Quantum mechanics*: Mosco con-

vergence of the Boolean Dirichlet forms delivers the canonical commutation relations from the gradient structure, and the Born rule follows from Gleason's theorem. *The Standard Model*: gauge bosons and the Higgs arise as inner fluctuations of the Dirac operator on $M \times F$, with $\phi_{\text{Higgs}} = \Pi(\psi)$: the modal condensate and the Higgs condensate are the same field at different scales. *Matter*: Boolean $\{0, 1\}$ distinctions are fermionic occupation numbers via the algebra of canonical anticommutation relations, making spin-statistics structural. Within the matter sector, the generation count $N_{\text{gen}} = 3$ is the natural structural reading: three Fano-plane sub-algebras of $\mathbb{O}$ through $\hat{\tau}$, read from Furey's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0, 6)$ via the chain-algebra $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$. Colour–generation distinguishability is settled in the negative inside PST (Computation 175): colour acts on the same three directions the generation count reads, and only its maximal torus preserves them (Computation 106).

The Standard-Model sector carries the theory's sharpest quantitative claim. The field-normalisation factor $Z^2 = \mathrm{e}^{-1}$ is derived within PST in the substrate limit (Computation 100), and its matching to the Standard-Model Higgs coupling is reduced to a field-strength-renormalisation identity whose internal renormalisation is complete (§12.5). What remains is one falsifiable premise, that the Standard Model is PST's effective theory below M_*, which a three-loop renormalisation-group test passes: $\lambda_{\text{SM}}(M_*) = 0.0930 \pm 0.0008$ against the predicted $\mathrm{e}^{-1}/4 = 0.09197$, a 1.4σ match with the residual set by the top-mass input rather than by PST (Computation 202). So $Z^2 = \mathrm{e}^{-1}$ is derived modulo (apart from) that premise: parameter-free, but not premise-free. **Predictions**

PST predicts a correction to the Casimir force that falls as the sixth power of the plate separation, with its coefficient derived; and it fixes that correction's size itself, far below what current experiments can see. The d^{-6} law is a departure from standard quantum field theory distinct in scaling from the d^{-7} retarded Casimir–Polder regime. The Gaussian convolution-kernel form is delivered structurally by the binomial-to-Gaussian central-limit-theorem limit of the substrate's Boolean projection kernel at matched scaling (Computation 73),

with calculable $-2/(3|D|)$ correction. The coefficient $\xi = 90/\pi^2 \approx 9.12$ is the matched-scaling limit, and it gives a diagnostic upper bound $d_0 \lesssim 2.3$ nm on the substrate coherence scale from current nanometre-Casimir data. That bound is set precisely because PST at $d_0 = 2.3$ nm predicts $\sim$0.19% at the well-measured $d = 160$ nm separation, the current precision floor at that scale. At smaller separations the signal scales as $(d_0/d)^2$: at the upper-bound d_0 the correction would be $\sim$1.9% at $d = 50$ nm and $\sim$12% at $d = 20$ nm. The theory's own canonical normalisation fixes d_0 nine orders lower, at $\sqrt{2}\,\hbar c/m_h = 2.2$ am (Computation 201), where the correction is 10^{-19} of the force: exact in form, far below present reach, with the d^{-6} exponent parameter-free and signature-distinguishable from all known material corrections.

The electroweak modal scale $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV is a parameter-free one-loop derivation. Two structural inputs fix it. The Landau–Ginzburg modal potential at threshold has quartic coefficient $\lambda_{\mathrm{PST}} = 1/4$ exactly (book §1.5), and the substrate boson–fermion mode count cancels by binomial identity ($N_B = N_F = 2^{D-1}$). The only remaining contribution to δm_h^2 is therefore the Higgs self-loop coefficient $C_H = 6\lambda_{\mathrm{PST}} = 3/2$, giving $m_h^2 = (3/2)M_*^2/(16\pi^2)$ and hence $M_* = 4\pi m_h\sqrt{2/3}$. The constant-function substrate vacuum mode is structurally distinct from the emergent Higgs doublet (Computations 93, 97). This is a scalar-sector result, a parameter-free one-loop derivation for the Higgs self-loop in isolation. Its extension to the full Standard Model (top, W, Z contributions) does *not* exist within PST (Computation 120). Three things stand in the way: the substrate cancellation is layer-disjoint from the Standard-Model loops (Computation 105); the Veltman residual at M_* is ≈ -3 with no zero crossing; and the bosonic top-partner sector that would cancel it is collider-disfavoured and absent from PST's spectrum (§12.2). M_* is therefore irreducibly a scalar-sector prediction, a definitive bound on the scope of the result, not a closure of it. Confronted with the Large Hadron Collider (LHC) in Computation 205, the scale survives in one reading only: a literal cutoff above which no field mode propagates, and a universal smearing of every vertex, are both excluded by the jet and dilepton spectra by large

margins, while the regulator on the Higgs doublet alone, which is what the derivation states, is consistent with every Higgs measurement and predicts a deficit of 23 percent in the Higgs transverse momentum spectrum above 650 GeV. That bin, measured at the Standard Model with 15 percent precision, would exclude the reading and with it M_* as a physical scale. The scalar sector also exhibits tree-level custodial SU(2) symmetry, which protects ΔT.

The cosmological-constant equation of state $w = -1$ follows conditionally from the substrate's pre-spatial framing via a steady-state-style non-dilution mechanism: the substrate's ongoing modal sublimation contributes a time-constant energy density that does not dilute under spatial expansion, so $\dot{\rho} = 0$ and the continuity equation $\dot{\rho} + 3H(\rho + p) = 0$ forces $w = p/\rho = -1$ exactly. The *magnitude* Λ_{obs}, the Yukawa hierarchies, and Cabibbo–Kobayashi–Maskawa mixing are contingent (they live in the realisation $T(C)$ rather than the structural P1–P3 layer) and not pinned by the postulates alone; the coherence length d_0 is not among them, being fixed at 2.2 am by the canonical normalisation (Computation 201). That normalisation, $c = 1$ for P3's kinetic term, which the value of M_* assumes, is a premise: Appendix C records it as assumed rather than derived, its rival reading excluded (§12.2).

Standing among substrate theories

Compared against the twenty-six other substrate-style & emergent-physics theories collected in Appendix D, PST is the only entry that does all five of the following at once.

(i) It posits a pre-geometric substrate, finite in internal structure rather than a finite universe: a Boolean distinction space, not a presupposed infinite-dimensional arena such as Wheeler superspace, a manifold, or a graph.

(ii) It derives spacetime as a limit theorem from that substrate, via Mosco convergence of the Boolean Dirichlet form to the Laplace–Beltrami operator on $\mathbb{R} \times S^3$, under the seven conditions A1–A7 of Chapter 18.

(iii) It derives the structural ingredients of the $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ gauge

group from the three postulates: the $\mathrm{Cl}(0,6)$ chain algebra delivers $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation, and $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ delivers an internal SU(2) on the emergent Dirac sector, with the full $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ following via Dixon's hyperspinor synthesis.

(iv) It places the three-generation count of matter in Furey's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0,6)$ as the natural reading.

(v) It contributes the derivation of a specific Standard-Model coupling value: $\lambda_{\mathrm{SM}}(M_*) = b \cdot \mathrm{e}^{-1} = 0.0920$, at 1.2% match. The substrate-side factor e^{-1} is derived from the foundational postulates (Computation 100). The Standard-Model-side matching is reduced *within* PST to a field-strength-renormalisation identity whose internal content is derived, resting on the emergence premise F4 (the Standard Model as PST's exact low-energy effective theory below M_*; §12.6.2.1) plus the renormalisation-group test, rather than inherited from the Chamseddine–Connes programme.

Other entries deliver one or two of these ingredients; PST delivers all five.

1. Introduction

Physics has two great theories, General Relativity for gravity and Quantum Mechanics for the very small, and they do not fit together. This book asks a different question from the usual one. Instead of trying to patch the two theories, it asks what the world would have to be made of for both of them to come out of it, and then starts from the smallest answer anyone could start from: that some things are simply not the same as other things.

1.1. The incompatibility of quantum mechanics and 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 quantise 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 quantising a nonlinear field theory; it is a symptom of the fact that both theories assume different things about the structure they presuppose [2, 112, 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 (QM) made compatible with special-relativistic spacetime. It presupposes more background structure than QM alone, namely a fixed

Minkowski metric [175] and causal order, and less than general relativity (GR), since it assumes no dynamical geometry. A foundational theory must derive all three; the natural demarcation is the emergence of the Lorentzian [181] 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 quantise it, or beginning with a causal order and asking how geometry emerges, PST begins at a level prior to both: a domain in which neither geometry nor causality is defined, and in which the only primitive is the capacity for distinctions to exist.

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

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

This domain, the precausal substrate, is not empty. It contains the asymmetric tension just described: 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 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.

1.5. The foundational object: the KO-dimension-ten triple

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 [185] to the Lorentzian, is a demand of consistency. It is the same demand that the incompatibility of quantum mechanics and general relativity, with which this chapter opened, leaves unmet by every background-assuming theory.

Connes' noncommutative geometry [48] answers that demand with a single integer, the real K-theory (KO) dimension, K for Grothendieck's *Klasse* and O for the orthogonal group, defined modulo eight by Bott periodicity [148, 149]. 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; Chapter 19). The Mosco limit [98] of the substrate's Dirichlet form produces a Lorentzian spacetime whose dimension the KO split fixes at four (below), 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 $4+6 = 10 \equiv 2 \pmod 8$. Ten is not selected by minimising

over totals: both summands are computed, the parity-selected substrate KO-6 (Computation 3) and the Mosco-limit spacetime KO-4 (Computation 4), so $4 + 6 = 10$ is arithmetic and the pair is unique (below). The foundational number of PST is therefore this total KO-dimension ten, split as the four of spacetime and the six of the internal factor, and it lands on two modulo eight, the Connes Lorentzian value. The split is conditional on the parity selection but requires no extra spectral-triple axiom beyond what Computations 3 and 4 verify; that the spacetime four is the unique mod-8 complement of the substrate six rather than an independent input is a finite arithmetic check (Computation 128).

PST uses *one* such extremal selection principle, applied twice. It was carried as three in the first edition, and it is now not carried as a commitment at all: Computation 148 reduced the three to one and Computation 168 derives that one from P1. Both steps are worth setting out, because a foundation that asks for nothing beyond its postulates is in better shape than one that asks for a principle beside them.

The principle is *maximality over the Hurwitz classification*: among the structures the four normed division algebras generate, take the largest. It has two applications.

- *Maximal directional structure*, for the substrate: $V_7 = \mathrm{Im}(\mathbb{O})$ is the largest vector part of any normed division algebra by the Hurwitz classification [164]. It is preferred to the 3-dimensional cross product because it is the maximal norm-multiplicative vector product (§1.7, §2.3).
- *Maximal-tensor hyperspinor envelope*, for the substrate–internal-algebra synthesis: $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ is the full Dixon tensor of all four Hurwitz division algebras, preferred to proper subalgebras such as Furey's [121] $\mathbb{C} \otimes \mathbb{O}$ alone (§13.7).

These are the same instruction evaluated on two constructions rather than two commitments: applied to vector parts it returns $\mathrm{Im}(\mathbb{O})$, and applied to tensor products it returns $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$, the two largest objects the classification

admits. That $\mathbb{R}$ is the tensor unit is why the envelope is a triple and not a quadruple product, which is a consequence rather than a choice.

The *minimal split* that the first edition listed as a third principle is not a selection at all. KO-dimension is additive, the substrate's KO-6 is verified (Computation 3) and the Mosco-limit spacetime's KO-4 is verified (Computation 4), so $4 + 6 = 10$ is arithmetic and the admissible pair is unique: there is nothing to minimise over. Had both summands been free, minimality would have selected $(0, 2)$, which the theory does not use (Computation 148). What the reduction does *not* do is remove the parity restriction that Computation 3 carries, since KO-6 holds on the odd-$|D|$ sub-limit. That conditioning is stated above and is unchanged; the weight moves from a named principle onto a restriction the theory already declares, rather than disappearing.

And the principle is not an extra commitment, because P1 already settles it. The first edition called it "natural given P1–P3 but not derived from them". That conceded too much, and the concession came from asking the wrong question. The question is not why the theory should prefer the *largest* structure, which would indeed need a principle. It is which structures can be named at all with the material P1 supplies. *Choosing a proper sub-object is not a smaller commitment than taking the whole; it is a larger one*, because one must also say *which* proper sub-object, and P1 supplies no selector.

That is measurable, and Computation 168 measures it. Nothing in P1 breaks the symmetry of the direction algebra, so the acting group is all of $\mathrm{Aut}(\mathbb{O}) = G_2$, and a choice costs nothing exactly when its G_2-orbit is a point. A unit imaginary direction costs six real parameters, its orbit being $S^6 = G_2/SU(3)$; a quaternionic triple costs eight, its orbit being $G_2/SO(4)$; every intermediate subspace costs something. Exactly two choices cost nothing: the zero subspace and the whole of V_7. The two clauses of P1 then settle which survives, the *sole*-primitive clause removing every intermediate choice and the *distinction-exists* clause removing the trivial one. The hyperspinor envelope goes the same way with a discrete rather than a continuous option set: of the sixteen

subsets of the Hurwitz classification, only the empty one and the whole one are specified without naming a member to omit, and the whole one is the triple product $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ of dimension 64, which recovers Computation 148's result rather than assuming it.

The criterion is not a synonym for "largest", and that matters. Rigidity is relative to the group: under $SU(3)$, the stabiliser of a single imaginary direction, a *one*-dimensional subspace is rigid, which is as far from maximal as a nonzero subspace reaches. The two criteria agree in PST only because P1 leaves the symmetry unbroken. So what looked like a preference for large structures is the absence of anything that could prefer a small one.

What this does *not* do should be as clear. It does not derive the Hurwitz classification, which is mathematics and is cited as such; it does not derive that the directional structure lies in that classification at all, which is norm-multiplicativity's work in §1.7 and §2.3, and is upstream of everything here; and it reads P1 as supplying no selector, which is how this book reads P1 throughout, but is a reading. A reader willing to add a datum may have whichever sub-object they like, and it will be their commitment rather than PST's. Within that scope the consequence is plain: the recurring "forced" phrasing in the body sections is to be read as forced, with no reading-down required.

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 (§13.7). The substrate's parity-selected real structure (Computation 3) delivers the Cl(0,6) chirality grading exactly (Computation 19: $\gamma_K = i\omega$ to machine precision in the Jordan–Wigner representation [105]), and Furey 2014 [121] obtains A_F from the algebra generated by the matter representation on the resulting minimal left ideal. The remainder of this book 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. It also isolates the two quantitative bounds in Chapter 19: the reduced bridge premise B behind $Z^2 = \mathrm{e}^{-1}$, and the negatively-settled full-Standard-Model extension of M_*.

1.5.1. The substrate–Higgs identification

The substrate's Landau–Ginzburg order parameter [8] ψ and the Standard-Model Higgs field ϕ [75] are not two analogous objects but a single field viewed at different energy scales: $\phi = \Pi(\psi)$, the modal condensate as seen from inside the instantiated geometry. The substrate-level Landau–Ginzburg (LG) sombrero is the same sombrero as the Standard Model (SM) Higgs potential, under the maps $\varepsilon \leftrightarrow \mu^2$, $b \leftrightarrow \lambda$ and $r_0 \leftrightarrow v/\sqrt{2}$ (doublet convention, $\lambda_{\mathrm{PST}} = \frac{1}{4}$; §1.7). The measured Higgs mass $m_h \approx 125$ GeV and the quartic are related under this map by $m_h = \sqrt{2\lambda_{\mathrm{SM}}}\, v$; the derived quartic is tested against the measured one at M_* in §12.2.

This one identification does a good deal of work later in the book, and it is worth saying now what that work is, so that the later chapters read as consequences rather than as fresh starts. At the electroweak scale the single complex ψ is not enough, and the threshold's own topology promotes it to the Higgs doublet, with ψ appearing as the neutral component H^0 in the broken phase (§13.4). The modal angle θ is then not an independent input but the

would-be Goldstone mode [163, 43] of that same condensate, eaten in the electroweak Higgs mechanism, with the photon emerging as the connection of the *unbroken* combination (§13.4). And because no long-range scalar degree of freedom survives in the infrared, the identification costs nothing observationally: $\gamma = \beta = 1$, so a substrate-wide condensate is compatible with the solar-system tests of gravity rather than in tension with them (§8.2). The scale at which the two viewpoints are matched, and the sense in which "different energy scales" is meant precisely, is §15.2; §17.5 sets the identification against the existing literature on electroweak symmetry breaking.

1.6. 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 artefact: 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. Above 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, 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 the Lorentzian spacetime $M = \mathbb{R} \times S^3$ directly, its dimension four fixed by the KO split; 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 follows under Dixon's hyperspinor

synthesis of the substrate's $\mathrm{Cl}(0,6)$ chain algebra and the emergent spacetime's $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ Dirac-spinor [23] algebra (Cartan classification of real Clifford algebras [153, 154]), both delivered by P1–P3 in conjunction with the Connes-spectral-triple framework. The principal pieces (§13.7) are three verified computations and one literature construction.

- Computation 3: the substrate's parity-selected real structure, on an odd number of distinction bits, delivers $\gamma_K^2 = +I$ with the Connes Standard-Model KO-6 sign triple $(\varepsilon, \varepsilon', \varepsilon'') = (+, +, -)$.
- Computation 19: γ_K equals the $\mathrm{Cl}(0,6)$ chirality grading $i\omega$ exactly in the Jordan–Wigner representation, so the chirality projector $f = (I + i\omega)/2$ comes directly from the substrate.
- Furey's 2014 chain-algebra construction [121], on the resulting minimal left ideal, delivers $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation of fermion content.
- Computation 66: the emergent Lorentzian Dirac sector $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ supplies the $\mathbb{H}$ factor of A_F by right-$\mathbb{H}$-multiplication on the spinor module (Dixon 2010 [146]).

The algebra identity holds, with no appeal to the Chamseddine–Connes–Marcolli classification [50]. The foundational object PST identifies is the product of M and F, a real spectral triple of total KO-dimension ten, whose $4 + 6 = 10$ minimal split compatible with a Lorentzian fermion action is derived above (§1.5).

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 quantising one and deforming the other. It is a matter of recognising 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.

1.7. The foundational postulates of PST

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 principal falsifiable prediction is a Casimir correction with d^{-6} scaling [11, 12, 13, 14, 15] and coefficient $\xi = 90/\pi^2$, with diagnostic amplitude ξd_0^2 bounded by nanometre Casimir data and fixed by the canonical normalisation at a value far below present reach (Chapter 15, Computation 201). The coefficient is parameter-free under the Gaussian projection-kernel choice, and the d^{-6} exponent is kernel-independent, derived in Computation 139.

PST rests on three foundational postulates (P1–P3). A fourth formal statement, the Laplace-type projection condition, is not postulated. It is reduced to P1–P3 via Mosco convergence of Boolean Dirichlet forms (§7.8) under A1–A7. The equidistribution condition A3 (the single named inequality $\eta_{K,R} < d_0^2/4$) is supplied by P3's maximum-entropy sublimation, so the Mosco limit closes under A1–A7 and equicoercivity A7 follows as a theorem given A1, A2, A4 and A3 (§18.1.6). A1, the matched-scaling condition, is architectural input rather than a substrate consequence, so the limit is conditional on it and is not a theorem of P1–P3 alone; A6's convergence-rate control is standard but not yet written out. The Laplace-type projection statement is given separately below for ease of reference. The postulates are stated in terms of a *direction algebra* V_7, the 7-dimensional vector space over $\mathbb{R}$ equipped with the octonion vector product.

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, (ii) $v\colon D \to V_7$ is the direction assignment, and (iii) δ is *faithful* to v: if

$v(a) \neq v(b)$ then $\delta(a, b)$, so that δ records every difference the substrate carries. Clause (iii) adds no structure to δ; it is a coherence condition between two components of the same triple, and without it δ would be *a* relation on D with substrate content beside it that δ cannot see, rather than *the* distinction relation this postulate names. It is a premise and is carried as one (§2.2, Computation 184). The dimension 7 of V_7 is not chosen: by Hurwitz's classification, the only normed division algebras over $\mathbb{R}$ are $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, $\mathbb{O}$, of total dimension $1, 2, 4, 8$; their *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, and we adopt this maximal direction algebra as the carrier of directional content. No geometric, causal, or temporal structure is presupposed; the directional content in V_7 is the only structural enrichment of bare distinguishability.

P2 (Asymmetric Vector Tension). The tension functional $T\colon \mathcal{P}(D) \to V_7$ takes its values in V_7. 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 (§14.2).

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

$$\mathcal{F}[\psi, \varepsilon] \;=\; \int_{\mathcal{C}} \big[-\varepsilon \, \langle \psi, \psi \rangle + \tfrac{1}{4} \langle \psi, \psi \rangle^2 + c \, |\nabla_{\mathcal{C}} \psi|^2_{V_7} \big] \, d\mu, \tag{1}$$

where $G_2 = \mathrm{Aut}(\mathbb{O})$ 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. We

adopt throughout the *doublet convention*: the quartic coefficient $\frac{1}{4}$ is the coefficient of the doublet bilinear $\langle\psi,\psi\rangle^2$, identified at the matched scale with the Standard-Model doublet form $V_{\mathrm{SM}} = \lambda_{\mathrm{SM}}(\Phi^\dagger\Phi)^2$ via $\langle\psi,\psi\rangle \leftrightarrow \Phi^\dagger\Phi$ (not the real-component norm; the two differ by a factor of two). In this convention $b = \lambda_{\mathrm{PST}} = \frac{1}{4}$ is the coefficient of $(\Phi^\dagger\Phi)^2$, the same convention as λ_{SM}, so that the field-normalisation ratio $Z^2 = \lambda_{\mathrm{SM}}(M_*)/b$ is a ratio of two same-convention couplings (§12.2, Computation 109). Under the radial reduction $\langle\psi,\psi\rangle \to \psi^2$ the postulate collapses to the scalar form $-\varepsilon\psi^2 + \frac{1}{4}\psi^4$ of equation (24). The gradient term is the Boolean Dirichlet form on $\{0,1\}^D$ extended pointwise on V_7 values:

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

(the sum over the cube of the per-configuration gradient of (21), up to the normalisation fixed there). 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).

Modal sublimation instantiates the maximum-entropy (μ-typical) configuration consistent with the tension data: among the substrate configurations that cross the threshold, the realised one is a typical configuration under μ, and in particular its emergence cloud is equidistributed at the closure rate that condition A3 states. This maximum-entropy selection is a typicality property of sublimation, internal to P3 and not a fourth postulate. Its role is to fix the realised emergence cloud as the equidistributed one, so that the equidistribution condition A3 of the Mosco limit (§18.1, §4.11) holds as a *consequence* of P3 rather than as an added premise. The same typicality applies to the content of T, the magnitude values and the elementary directions that the tension data consist of. The realised content is a typical point of the content space under its invariant measure, and a typical point lies off every closed set of

positive codimension. Two such sets matter: no two magnitudes coincide off the distinction-free configurations, which is the Born condition on the law of T, and the elementary directions do not sum to zero, which is non-balance (Computations 193, 195). Both are therefore supplied on the same footing as A3. This is a reading of the clause, not a fourth postulate.

1.7.1. Derived statement: the Laplace-type projection

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

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

so that $\widetilde{K}(x, y; t) = \langle x|e^{-t\Pi(L)}|y\rangle$. This statement is *not* postulated: it follows from P1–P3 by Mosco convergence of the V_7-valued Boolean Dirichlet form to the standard Laplace–Beltrami form [169] on $M = \mathbb{R} \times S^3$ along the selected direction $\hat{\tau}$ (§7.8). Structurally PST rests on three postulates only, and most of P2's operative content reduces to P1 with P3's typicality (Chapter 3), so the independent commitments are fewer still.

The complex-scalar Landau–Ginzburg formulation is 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, §14.2).

The Bernoulli measure [168] $\mu = \bigotimes_{a\in D} \mathrm{Bern}(\frac{1}{2})$ is not a postulate; it is the unique measure on $\mathcal{P}(D)$ invariant under relabelling and under per-property complementation (§2.5). The Lorentzian signature $(+, -, -, -)$ is

derived from hyperbolic well-posedness of the Goldstone wave equation along $\hat{\tau}$ (Computation 2). See Chapter 19 for the full accounting and for the open research that remains.

1.8. Structure of this book

The remainder of this book develops PST in full. Chapter 2 develops postulate P1 (property differentiation) with its mathematical expression. Chapter 3 develops postulate P2 (asymmetric vector tension). Chapter 4 develops postulate P3 (modal sublimation): the modal threshold and its variational formulation in the first half, and the sublimation transition formalised as a categorical functor with explicit construction in the second half. Chapter 5 addresses the coinstantiation of spacetime and causality. Chapter 6 consolidates the emergence chain. Chapter 7 derives energy, matter, and fields as projections of asymmetric tension, on four analytic results stated there, three proved in Chapter 18 and the fourth in §8.10, equicoercivity A7 among them given A1, A2, A4 and A3. Its equidistribution condition A3 (the closure margin $\eta_{K,R} < d_0^2/4$ below which the continuum manifold limit closes, for the emergence-determined realisation ρ_n) is supplied by P3's maximum-entropy sublimation (§18.1.6). Chapter 8 checks that the projected geometry is Einstein's: spacetime curvature, the weak-field limit that recovers Newtonian gravity, horizons, Lovelock uniqueness, the equivalence principle, and Newton's constant [174] as a unit conversion fixed by the projection scale rather than a free parameter. Chapter 9 works out what the substrate leaves behind when nothing is there: the sign of the cosmological constant, the degenerate vacuum manifold, and angular momentum as its Noether charge. Chapter 10 recovers quantum mechanics: the Born rule from the Bernoulli measure, the measurement postulate from threshold-crossing, the canonical experiments, composite systems and no-signalling, and the canonical commutation relations. Chapter 11 identifies the fermionic sector through the canonical anticommutation relation (CAR) algebra and derives the statistics of the fermionic sector from Boolean exclusion (exclusion forces the CAR and composition parity fixes the statistics; the spin label is attached by the spinor bundle). Chap-

ter 12 builds the projected quantum fields and derives the Higgs quartic, the number the book puts against measurement. Chapter 13 derives the Standard Model gauge group from the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ (closing the Conjecture-1 algebra identity, §13.7), fixes the hypercharge assignment, and establishes gauge anomaly cancellation as a structural consistency check. Chapter 14 derives the three-generation count from the directional content of V_7, states the structural-scope theorem that places the Yukawa couplings and the Cabibbo–Kobayashi–Maskawa (CKM) matrix in the configuration's content, and obtains the chirality of $\mathrm{SU}(2)_L$ as a structural consequence. Chapter 15 identifies the discreteness scale d_0 with the Landau–Ginzburg coherence length of the modal vacuum, and develops the Casimir correction, the d^{-6} scaling a necessary signature and the parameter-free coefficient $\xi = 90/\pi^2$ fixing the diagnostic amplitude ξd_0^2, exact in form and far below present reach. Chapter 16 develops PST's cosmology: the Friedmann [68] equations are derived, the Cosmological Principle is established as a theorem of the Bernoulli measure, and ongoing instantiation is shown to deliver the equation of state $w = -1$ structurally, conditional on the substrate's pre-spatial framing (§9.1). Chapter 17 positions PST in relation to existing frameworks. Chapter 18 carries the hardest mathematics, the Mosco convergence of the projected geometry and the Dirac-operator convergence under assumptions A1 to A7, which the rest of the book needs and no reader needs to follow. Chapter 19 catalogues the structural and mathematical research problems that the theory resolves, sector by sector. Chapter 20 concludes. Appendix A is a consolidated table of the symbols recurring through the book, grouped by sector, each with the section where it is first introduced. Appendix B is the register of the computation scripts cited throughout as Computation N; Appendix C consolidates the status of every substantive claim of the book; Appendix D tabulates the comparison to other substrate and emergent-physics theories; Appendix E collects the bibliographic references.

1.9. 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 book 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. Property Differentiation: The Sole Primitive

Before there is space, or time, or matter, there has to be at least one difference. This chapter says what a difference is when you are not allowed to use anything else to explain it. Two things differ simply because they are not the same thing. Nothing more is needed, and nothing less would do.

2.1. Differentiation by non-identity alone

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.

2.2. The substrate triple (D, δ, v)

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

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

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

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

- $v\colon D \to V_7$ is the direction assignment: each property $a \in D$ has inherent directional content $v(a) \in V_7$ in the 7-dimensional direction algebra.
- δ is *faithful* to v: a difference in directional content is recorded as a distinction,

$$v(a) \neq v(b) \implies \delta(a, b) \tag{6}$$

One condition beyond (4) *and* (5) *is imposed on* δ*, and it is a condition of*

coherence with v rather than of structure. Faithfulness (6) adds nothing to δ that is not already in the triple: it requires only that the relation whose job is to record distinctions record the ones the triple carries. Beyond it nothing is imposed. There is no transitivity requirement, no totality requirement, no metric on D, no topology, and no ordering.

Faithfulness is a premise and is carried as one (Appendix C). It is not a theorem of (4) and (5): the two-property substrate with δ empty and $v(a) \neq v(b)$ satisfies both and violates it (Computation 184). Its motivation is the one just given, that without it δ is not *the* distinction relation this chapter names; *its consequence for Born agreement was found afterwards*, in Computation 184, and is recorded in that order deliberately. What it buys structurally is set out there and is worth stating here, because a reader will ask: contraposed, (6) says that two properties carrying no distinction between them share their directional content, so δ contains the complete multipartite graph on the partition of D by v-value and the only freedom left lies inside a v-block. Where v is injective, δ is complete. 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.

2.3. The 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 $\times : V_7 \times V_7 \to V_7$ 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 $\mathbb{R}$.

2.4. Configurations and configuration space

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

$$v(C) \;=\; \sum_{a \in C} v(a) \;\in\; V_7. \tag{7}$$

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 Chapter 3.

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.

2.5. The canonical measure on configurations

Before assembling the functional (Chapter 4), the measure μ on $\mathcal{C}$ must be fixed. The variational formulation demands integration over all configurations, and the choice of measure is not arbitrary: it must follow from the structure of (D, δ) alone, without invoking any geometric, probabilistic, or metric primitive that would introduce a new assumption at the level of the substrate.

The configuration space $\mathcal{C} = \mathcal{P}(D)$ is the power set of D: the collection of all subsets of the property domain. The structure (D, δ) carries a natural symmetry group, its automorphism group $\mathrm{Aut}(D, \delta)$, consisting of all bijections $f : D \to D$ that preserve the distinction relation: $\delta(a, b) \iff \delta(f(a), f(b))$ for all $a, b \in D$. Every such automorphism lifts canonically to a bijection on $\mathcal{C}$ by acting elementwise: $f \cdot C = \{f(a) : a \in C\}$. Since the only primitive is

(D, δ), and since the definition of a configuration is simply a subset of D, the measure μ must be invariant under this lifted action: no configuration should receive more weight than any automorphically equivalent configuration.

The key observation is that in the precausal substrate, all properties in D are structurally equivalent: the distinction relation δ is irreflexive and symmetric but otherwise unconstrained, so there is no intrinsic difference between any one property and any other beyond their bare identity. No property is marked, distinguished, or preferred. This means that the full symmetric group on D is contained in $\mathrm{Aut}(D, \delta)$ as a subgroup: any permutation of D is an automorphism, since δ treats all pairs equally, being complete on the properties of distinct v-value that faithfulness distinguishes (below). 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 canonical measure on $\mathcal{P}(D)$ is the uniform measure, assigning equal weight to all 2^n subsets:

$$\mu(C) = 2^{-n} \qquad \text{for all } C \in \mathcal{C} \tag{8}$$

Permutation invariance alone does not force this: any weighting that depends only on the cardinality $|C|$ is permutation-invariant. What forces it is the per-property form of complementation invariance. The primitive (D, δ) draws no distinction between the presence and the absence of each *individual* property, so for every $a \in D$ the single-property flip $\sigma_a \colon C \mapsto C \triangle \{a\}$ is a symmetry of the substrate. The group generated by the flips acts transitively on $\mathcal{P}(D)$, so the unique flip-invariant probability measure is the uniform one; permutation invariance then holds automatically.

For an infinite domain D, the natural extension is the Bernoulli product measure $\mu = \bigotimes_{a \in D} \mathrm{Bern}(1/2)$, the unique probability measure on $\{0,1\}^D$ invariant under all permutations of the index set D and under every single-property flip σ_a. The flips act transitively on every finite cylinder, forcing

the uniform marginal on each finite window, and Kolmogorov extension then gives the product measure:

$$\mu\big(\{C \in \mathcal{C} : C_0 \subseteq C\}\big) = 2^{-|C_0|} \qquad \text{for all finite } C_0 \subseteq D \tag{9}$$

This is the projective limit of the uniform measures on all finite sub-domains of 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.

2.5.1. 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 every single-property complementation $\sigma_a \colon C \mapsto C \triangle \{a\}$, $a \in D$ (the global complementation $C \mapsto \bar{C}$ is their composition), is the Bernoulli product measure $\mu = \bigotimes_{a \in D} \mathrm{Bern}(1/2)$.

The proof is immediate from (ii) alone: the group generated by the single-property flips acts transitively on $\mathcal{P}(D)$ for finite D, and on every finite cylinder for infinite D, so an invariant measure weights all configurations (all cylinders over a given window) equally, which is the Bernoulli product measure; (i) then holds automatically. Condition (ii) is the exact measure-theoretic expression

of P1's symmetry between the presence and the absence of each distinction. The per-property form is essential: global complementation together with permutation invariance would not suffice (a mixture supported on $\{\emptyset, D\}$ satisfies both), which is why the flip invariance is the honest hypothesis.

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.

2.6. Summary

This chapter develops P1, the property differentiation primitive, and assumes only P1. No space, no time, no energy, no forces, and no laws of nature are assumed here. PST's sole ontological primitive is the capacity for distinctions to exist; from this alone, the substrate's internal structure (the set D of primitive properties and the Bernoulli measure μ) is derived. The Bernoulli measure $\mu = \bigotimes_a \mathrm{Bern}(\frac{1}{2})$ is derived as the unique measure on $\mathcal{P}(D)$ invariant under relabelling and per-property complementation; adopting the direction algebra $V_7 = \mathrm{Im}(\mathbb{O})$ as the maximal directional structure (§1.5) was carried as a modelling choice until Computation 168, and is now a consequence of P1: a proper subalgebra such as the cross product V_3 on $\mathbb{R}^3$ costs a selector P1 does not supply, and V_7 costs none. What remains upstream is the norm-multiplicativity that places the substrate in the Hurwitz classification at all (§1.7, §2.3). No conjectures enter at the postulate level. The two further postulates, P2 (asymmetric vector tension, Chapter 3) and P3 (modal sublimation, Chapter 4), supply the substrate's asymmetry and its threshold-crossing transition respectively.

3. Asymmetric Vector Tension: The Preferred Direction

Think of a tug of war. If the pulls on the rope balance, nothing moves; if they do not, the rope goes one way. This chapter gives the substrate its own version of a pull, and the whole book depends on that pull failing to cancel. What is left over when it fails to cancel is what everything else gets built from.

This chapter develops P2, the asymmetric vector tension postulate. Where P1 (Chapter 2) supplies the substrate's primitive distinguishability, P2 supplies its asymmetry: the structural fact that property differentiation carries a directional imbalance with V_7-valued tension. P2 is a separate postulate, and the precise sense in which it is separate is worth stating rather than waving at: it is not derivable from P1 *alone*, but most of its operative content is. Computation 127 reduces it to P1 together with magnitude positivity (14), leaving one non-degeneracy condition, the non-balance of the elementary directions, which is logically independent of both. So what P2 adds beyond P1 is narrower than a whole postulate's worth, and that residue is supplied by P3's typicality clause (Computations 193, 195; Appendix C), so nothing is left standing on the structural front.

3.1. The structural origin of 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 opposite tension, $T(\bar{C}) \neq -T(C)$. This structural

non-equivalence between a configuration and its complement is what gives asymmetric tension its name and its content. It is not that some properties are heavier or more energetic than others. It is that the configuration, considered as a whole, stands in a structurally non-equivalent relation to its complement: their mutual imbalances do not cancel.

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

3.2. The tension functional and the direction algebra V_7

This modal imbalance is represented by the tension functional T, a V_7-valued map on configurations:

$$T\colon \mathcal{C} \to V_7. \tag{10}$$

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 $v\colon D \to V_7$ 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 total preorder on configurations (any two configurations can be compared by $|T|$, with distinct configurations permitted to share a value), and this ordering is what the variational analysis of Chapter 4 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 (§14.2). The V_7-valued representation is forced by the Hurwitz constraint inside P1: a real-scalar representation would discard the directional information $v(a)$ encoded in the elementary properties of D.

3.3. Epistemological status of the formalism

A word on the epistemological status of the formalism introduced in this and the following chapters 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 artefacts 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.

3.4. Complement asymmetry and the substrate's directional bias $\hat{\tau}$

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

$$T(\bar{C}) \neq -T(C) \tag{11}$$

generically. Equivalently, the total substrate tension, the tension $T(D)$ of the maximal configuration $C = D$,

$$T(D) \neq 0, \tag{12}$$

does not vanish *generically*: D contains distinct pairs, so its edge-count magnitude is maximal, but the net V_7-vector $T(D)$ is nonzero precisely when the elementary directions do not antipodally cancel, $\left|\sum_{a\in D} v(a)\right| > 0$. This non-balance is logically independent of magnitude positivity (14) and is the genuine content of complement-asymmetry (16) (both stated as structural constraints on T in §4.1 below); it is not forced by (D, δ) alone, and its standing is a stability statement rather than a probabilistic one: the balanced locus is closed with empty interior, so non-balance is an *open* condition whose distance to failure is exactly $\left|\sum_a v(a)\right|/\sqrt{|D|}$ (Computation 127, Computation 152). It is *adopted* rather than derived: Computation 169 sharpens it considerably, showing that balance would need every property's conditional law to carry an atom on the cancelling value and so requires total coordination, but the step that would have turned that sharpening into a derivation is withdrawn (§3.5). This non-vanishing of the substrate's total tension is the operative content of P2: $T(D)$ is the substrate's net directional bias in V_7, a configuration-independent V_7-vector fixed by D itself rather than by any partition into a configuration and its complement, and it is what supplies the preferred direction $\hat{\tau}$ for modal sublimation. It is the precise substrate-level form of the complement-asymmetry introduced earlier in this chapter, and is logically distinct from the configuration-level inequality $T(C) \neq T(\bar{C})$, which can hold

or fail for individual configurations independently of whether the net bias vanishes.

What P2 supplies is the *directional* character of this imbalance: that the modal tension is V_7-valued, carrying a net direction, rather than a bare scalar magnitude. Given that directional character, the net direction $\hat{\tau} = T(D)/|T(D)|$ is well defined and nonzero, since $|T(D)| > 0$ under non-balance, and it is this direction that modal sublimation locks into (P3). The net imbalance is therefore generically a nonzero V_7-vector:

$$\Delta T(C) \;:=\; T(C) - T(\bar{C}) \;\neq\; 0 \quad \text{in } V_7, \text{ generically in } C \tag{13}$$

(as with (11), the balanced locus for any fixed C is measure zero but not empty).

3.5. The modal and ontological status of $T(C)$

It bears emphasis that $T(C)$ is not energy and $\Delta T(C)$ is not energy difference. The tension functional assigns a V_7-valued imbalance 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. What carries a real-number value is the *magnitude* $|T(C)| \in \mathbb{R}^+$; its use to represent the degree of incoherence is a mathematical convenience, since the reals form the simplest ordered field with the topological properties needed for the variational analysis of Chapter 4, and the ordering on $\mathbb{R}^+$ provides a natural notion of more or less tension without requiring any further geometric or physical interpretation. The values taken by $|T(C)|$ 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 chapter 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.

A candid qualification is in order. Since δ is symmetric, the non-identity relations themselves carry no directionality; what the theory asserts is the complement-asymmetry condition $T(\bar{C}) \neq -T(C)$ of (11), 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. The status of this asymmetry is now located precisely: it reduces to P1 together with magnitude positivity (14) except for one non-degeneracy condition, the non-balance of the elementary directions $\left|\sum_{a\in D} v(a)\right| > 0$, which is the genuine irreducible content of (16). Non-balance is not forced by the structure of D. What can be said without introducing a measure is stronger than a measure-zero remark, and borrows nothing that a measure would have to supply.

Non-balance is an open condition. The balanced locus $B = \{v : \sum_a v(a) = 0\}$ is closed and of codimension seven in the space of direction assignments: a linear subspace in the unrestricted model, and for $|D| \geq 3$ a compact real-algebraic set of the same codimension at every balanced assignment

except the all-parallel ones, where it drops to six (Computation 152); at $|D| = 2$ balance forces the antipodal configuration, so the codimension is six throughout. In every case B is closed with empty interior, and its complement is open and dense, which is all the argument requires. Non-balance holds on a whole neighbourhood of any assignment that has it and survives every sufficiently small perturbation, while balance survives none: the distance from an assignment to exact balance is precisely $\left|\sum_a v(a)\right|/\sqrt{|D|}$. This is a statement about stability rather than probability, it needs no measure over v, and the measure-zero claim is its weaker corollary, true under every absolutely continuous law.

Nor does symmetry favour balance. The natural objection is that a substrate with no structure to break should be maximally symmetric, and that maximal symmetry ought to mean cancellation. It means the opposite. The assignments left fixed by the substrate's own relabelling group are exactly the constant ones, and those are maximally *un*balanced, $\left|\sum_a v(a)\right| = |D|\,|v|$ (Computation 152). Exact cancellation requires a pairing of the properties, a fixed-point-free involution on D, structure over and above (D, δ), which is precisely the coordination the substrate does not supply.

What is still not forced. Symmetry alone does not deliver non-balance, and the point should be conceded sharply rather than left vague. Let each $v(a) = \epsilon_a u$ for one shared u and independent signs $\epsilon_a = \pm 1$. That law is invariant under the diagonal G_2 action, under every relabelling of D, and under per-property complementation $v(a) \mapsto -v(a)$, the exact analogue of the flips that pin the Bernoulli measure in §2.5, yet it balances with probability $\binom{|D|}{|D|/2}2^{-|D|}$, about 0.31 at $|D| = 6$ (Computation 152). The symmetry package PST has articulated is therefore demonstrably insufficient, and the residue is a coordination assumption, not a symmetry one. What suffices is much weaker than independence: balance can carry positive probability only if the direction of *every* property is pinned by the rest to one exact value with positive weight, so it is enough that a single property, anywhere in D, is not so pinned. Whether that last premise is itself derivable, or is the irreducible primitive,

was the sharpened open question.

Non-balance is one face of one condition (Computation 193): the contingent content of T is a magnitude and a direction, the magnitude must avoid the tie set (closed of codimension one; that is the Born condition on T, §10.3.2.5) and the direction must avoid the balance set (closed of codimension seven), each face open and dense, null under continuous content and of positive mass under discrete content, the two independent of each other and the theory fixing neither. The argument that would close both faces is the maximum-entropy one (Computation 195): the maximum-entropy law on the content, relative to the invariant reference measure, has a density, so both exceptional sets are null under it and $\mu_{\pm}$ below is its antipode rather than its rival; what it needs is that P3's typicality clause reach the content of T as well as the configuration, which is the reading this book makes (§1.7): non-balance is supplied on A3's footing and has left the structural front with the Born condition. What follows records the direction face's own structure, which is what the supplied condition rests on.

An argument that it is derivable was advanced and is withdrawn. Computation 169 sharpens the condition considerably but does not close it, and the sharpening is worth having on its own terms. Balance requires $v(a) = -\sum_{b \neq a} v(b)$ for each a: one point of S^6, fixed by the others. So Pr[balance] is exactly the expected conditional mass sitting on that point, and if the conditional law of even *one* property is atomless the whole probability is zero (Computation 169). Balance therefore needs every property's conditional law to carry an *atom*, placed exactly on the cancelling value. So balance requires *total coordination*: every property pinned by the rest, and one unpinned property anywhere in D destroys it.

The step that would have closed it does not work, and the counterexample is already in this book. The argument advanced was that an atom is a distinguished direction, costing six real parameters under the group P1 leaves unbroken, as Computation 168 measures for the direction algebra. That prices

an *unconditional* atom at a *fixed* point of S^6, while what the chain above delivers is a *conditional* atom whose location depends on the other properties, which is a correlation and names no direction at all. Computation 152's own $\mu_\pm$ is the refutation and Computation 169 measures it: with $v(a) = \epsilon_a u$ and u drawn *uniformly*, the law is rotation-invariant, names no direction, pays nothing, has *maximal* conditional atoms, and balances 31% of the time. *So non-balance remains a condition*, and what would close it is an argument about correlation rather than about direction: a reason the substrate cannot carry a shared hidden variable of $\mu_\pm$'s kind. Computation 168's selector argument is not that argument.

The 31% *is forced rather than chosen, and it does not scale* (Computation 178). The closed form is already stated above, so what is added is not the formula but the fact that the symmetry package *leaves no other law to choose*, which turns one counterexample into a bound on the whole class. Per-property complementation $v(a) \mapsto -v(a)$ is, on the all-parallel family, exactly $\epsilon_a \mapsto -\epsilon_a$, and the single-coordinate flips act *simply transitively* on $\{\pm 1\}^{|D|}$, so an invariant law on the sign vector can only be the uniform one. $\mu_\pm$'s independent signs are therefore not a modelling choice but the sole complementation-invariant option, and the balance probability is the forced value $\binom{|D|}{|D|/2}2^{-|D|}$, which is $5/16$ at $|D| = 6$ and behaves as $\sqrt{2/\pi|D|}$: some 10^{-40} at a substrate cardinality of 10^{80}, and *exactly zero at every odd* $|D|$, a sum of an odd number of ± 1 being odd. Independently of that family, complementation invariance alone caps *any* law at $\Pr[\text{balance}] \leq 1/2$, since flipping a single property sends $\sum_a v(a)$ to $-2v(a) \neq 0$ and so pairs each balanced configuration with an unbalanced one of equal measure; the bound is attained at $|D| = 2$. *None of this derives the condition*, and it is not offered as doing so: a probability bound is a statement about a law and not about the one configuration the substrate has, the vanishing is established for the family the counterexample inhabits, and $1/2$ forbids certainty rather than balance. What it establishes is that the obstruction does not grow with the substrate, and that the symmetry package already asserted is what prevents it.

The measurement makes the reach of this concrete. Take the counterexample above at $|D| = 6$, which balances 31% of the time, and leave a *single* property unpinned while the other five stay maximally coordinated: balance does not become rare, it becomes null, and $\Pr[|\sum_a v(a)| < \epsilon]$ falls as ϵ^6, the exponent of the all-parallel stratum that Computation 152 identified independently. Five coordinated properties cannot rescue balance from one that is not.

Scope, which is the honest part. The sharpening above is about the continuum of directions P1 supplies, and not about the quantised variant where the substrate is handed fourteen distinguished axes at the outset; there the premise fails by construction, and Computation 152 measures balance at 1/14 for $|D| = 2$. That variant is a different theory (§19.6). The counterexample above and that scope note are *two* mechanisms rather than one, which is worth separating: the quantised variant balances by confining the directions to a distinguished finite set, so the sum lands in a lattice containing zero, whereas $\mu_\pm$ confines nothing and fixes no axis, and balances because the properties are *correlated* (Computation 169).

It is worth stating plainly how much rests on this condition, because it is more than the preceding paragraph suggests. Non-balance carries two of this book's principal claims, and it carries them independently of one another.

The first is that anything is instantiated at all. With $\sum_{a \in D} v(a) = 0$ the net direction $\hat{\tau} = T(D)/|T(D)|$ is not small but undefined, the excess $\varepsilon(C) = |T(C)| - \tau$ never becomes positive, and modal sublimation does not occur: the instantiated branch is empty. The second is the claim of §4.1, that the angular direction of the vacuum is pinned by the substrate rather than spontaneously selected. That claim reads $\hat{\psi}$ off $\hat{\tau}$; with $\hat{\tau}$ undefined there is nothing to lock to, and the S^6 of equally deep minima reverts to exactly the spontaneous-arbitrary level that PST otherwise improves upon.

The answer to why there is anything and the answer to why this vacuum therefore rest on the same non-degeneracy condition, supplied on A3's footing by P3's typicality (§1.7) and carried in Appendix C. Nothing is concealed by

this: the condition is stated and its status given above. But it is a single point of failure, and it is better named here than left for a reader to assemble.

The genericity defence just given can be strengthened, and it is worth separating what it establishes from what it assumes. That the balanced locus is measure zero is a statement about a measure, and PST fixes none over $v\colon D \to V_7$; the defence therefore borrows what the same sentence concedes is absent. The contrast with this book's own practice elsewhere is the sharper point: the Bernoulli measure on $\mathcal{P}(D)$ is not adopted but derived, as the unique measure invariant under relabelling and under per-property complementation (§2.5). Uniqueness under an invariance is the form of argument PST uses when it can.

A defence of that shape is available here without any measure at all. Exact cancellation is a single global constraint relating the values of v across the whole of D. For it to hold exactly, something must coordinate those values, and at the substrate level there is by construction nothing that could: no prior structure, no dynamics, and no law relating one elementary direction to another. P1 grants only that properties may differ. Balance is therefore not improbable but unmotivated, a constraint with nothing available to impose it, which is the same move this book makes against nothing: that nothing has no way to be, rather than that nothing is unlikely.

The natural objection, that nothing prevents balance either, is worth meeting directly. The asymmetry is that balance is a condition, an equation required to come out exactly right, whereas non-balance is the absence of one. Conditions require grounds; absences do not. A theory is entitled to decline to explain why a constraint it never imposed fails to hold.

This replaces a probabilistic defence with a structural one, and the two turn out to be the same argument rather than rivals. A measure over v factors into a directional part on S^6 and a radial part, and neither is what is actually being assumed. G_2 acts transitively on S^6, so that sphere carries a *unique* G_2-invariant probability measure: the directional part is fixed by uniqueness

under an invariance, exactly as the Bernoulli measure on $\mathcal{P}(D)$ is. And the radial part turns out not to matter, balance having probability zero under every radial law that grants each property some directional content at all (Computation 141). What carries the conclusion is that nothing coordinates the properties: independence across properties is the strongest form of that, and the weaker form that suffices is the one stated above, that no property's direction is pinned by the rest. Either is precisely the claim argued for just above, that nothing at substrate level is available to coordinate the values. The structural argument supplies the probabilistic one's missing premise.

This does not derive non-balance. It sharpens what P3's typicality is asked to supply: no longer a measure over v, but the absence of coordination over v, which is where that typicality now carries it (§1.7). 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.

3.6. Philosophical lineage: Leibniz, Spencer-Brown, Hegel

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* [61] 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 [67] anticipates the structure in dialectical form: pure being has no determinate content and resolves into nothing; the synthesis is becoming, the first concrete category. The precausal substrate has no geometric or causal content and cannot remain in that indeterminate state once tension crosses τ: a modal sublimation that parallels Hegel's becoming while supplying a precise mathematical mechanism in its place. Leibniz grounded explanation in the logical necessity of reason; Spencer-Brown grounded logic in the primitive of distinction; PST grounds geometry in the logical necessity of tension. What PST adds that none of its predecessors supplies is the transition: a precise variational mechanism by which the primordial asymmetry accumulates to a critical value and generates, through a well-defined phase transition, the geometric structures that all subsequent physical law presupposes.

3.7. Summary

P2 supplies the asymmetric vector tension: each substrate configuration carries a V_7-valued imbalance $T(C)$, structurally distinct from P1's primitive distinguishability and from P3's threshold-crossing transition. P2 is the second of PST's three foundational postulates; together with P1 (Chapter 2) it constitutes the ontological layer of the theory, with P3 (Chapter 4) supplying the dynamical layer.

The chapter's only postulate is P2 (asymmetric vector tension, V_7-valued). As in the previous chapter, taking $V_7 = \mathrm{Im}(\mathbb{O})$ as the direction algebra (§1.5) is no longer a modelling choice: a proper sub-algebra such as V_3 costs a selector P1 does not supply, and V_7 costs none (Computation 168). What is derived is the nonvanishing of $T(D)$ *given* the non-balance it rests on; the condition itself is not. It is not forced by (D, δ) alone, and what it has in its favour is stability rather than derivation: it is an *open* condition, the balanced locus being closed with empty interior so that non-balance survives every small perturbation with distance to failure exactly $\left|\sum_a v(a)\right|/\sqrt{|D|}$ (Computation 127, Computation 152). Computation 169 sharpens it further: balance requires every property's conditional law to carry an atom on the

cancelling value, so *total* coordination is needed and one unpinned property anywhere destroys it. *It does not close it*, and an argument that it did is withdrawn (§3.5): the coordination is a correlation, not a named direction, and Computation 152's invariant $\mu_{\pm}$ exhibits it at zero cost while balancing. So non-balance is supplied by P3's typicality rather than derived from the symmetry package; no conjectures enter at the postulate level.

4. Modal Sublimation: How the Possible Becomes Real

Water does not get steadily more solid as it cools. At zero degrees it becomes ice, all at once. This chapter argues the substrate does something similar: when the leftover pull passes a certain size, a configuration stops being merely possible and becomes real. That sudden change is what the book calls modal sublimation.

This chapter develops P3, modal sublimation. Where P1 supplies primitive distinguishability and P2 supplies its asymmetric tension, P3 supplies the transition mechanism: a Landau–Ginzburg threshold beyond which the uninstantiated substrate cannot maintain itself, producing instantiated geometry along the substrate's net direction. The chapter has two parts: the threshold mechanism that drives the transition (§4.1 and following) and the sublimation transition itself, formalised as a categorical functor and constructed explicitly.

4.1. 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. (This boundary is *modal*: a logical and structural bound on uninstantiated imbalance, the modal threshold τ. It is not a spatial or cosmological bound on the size or finitude of the universe. That separate question is addressed in §1.7 via the emergent topology $M = \mathbb{R} \times S^3$.)

To understand why this boundary must exist, recall the character of the precausal substrate. It is a domain that is defined entirely by what can be coherently maintained without spacetime, without causal structure, and without any form of realisation. A configuration exists in the substrate insofar as it is a coherent arrangement of distinctions: its properties are differentiated from one another, and that differentiation constitutes its identity. Coherence

here is not a quantitative threshold in the ordinary sense. It is a modal condition: a configuration is coherent as an uninstantiated structure if and only if it can exist as such without contradiction.

Now, asymmetric tension is the imbalance inherent in any differentiated configuration. As the number and depth of distinctions within a configuration increase, 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:

$$|T(C)| \;\geq\; 0, \quad |T(C)| = 0 \;\Leftrightarrow\; C \text{ has no distinct pairs} \tag{14}$$

$$|T(C \cup \{a\})| \;\geq\; |T(C)| \quad \forall\, a \notin C \tag{15}$$

$$T(\bar{C}) \;\neq\; -T(C) \quad \text{generically} \tag{16}$$

Equation (14) is magnitude positivity, equation (15) is magnitude monotonicity, and equation (16) is complement vector asymmetry. The last says that the substrate's total tension $T(D)$, the tension of the maximal configuration $C = D$, is *generically* a nonzero V_7-vector, providing the substrate-level directional bias $\hat{\tau}$ that P2 names. Generically rather than necessarily: the two constraints do not establish the same thing. (14) forces the *magnitude* $|T(D)| > 0$, since D contains distinct pairs, whereas the *V_7-vector* $T(D)$ is nonzero only where the elementary directions fail to cancel antipodally. That non-balance is independent of (14), and it is *unforced*: an open condition with a sharp margin, a stability statement rather than a probabilistic one (§3.4, Computation 152). The magnitude constraints (14)–(15) and the directional constraint (16) bear on different aspects of T, its scalar magnitude and its V_7-direction. A minimal example for the *magnitude* structure is the induced edge count:

$$T_{\text{ex}}(C) \;=\; \big|\big\{(a,b) \;:\; a,b \in C,\; a \neq b,\; \delta(a,b)\big\}\big| \tag{17}$$

Equation (14) holds for any configuration with at least two distinct properties. Equation (15) holds by inspection: adding property a to C increases the count by $|\{b \in C : \delta(a, b)\}|$. At the scalar (magnitude) level, the analogue of (16) (that $T_{\text{ex}}(C) \neq T_{\text{ex}}(\bar{C})$) holds whenever the induced subgraphs on C and $\bar{C}$ have different edge counts, generic for non-regular distinction structures. Equations (14)–(16) 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 scalar example fixes the *magnitude* but carries no V_7-direction; the directional content comes from aggregating the elementary directions $v(a)$ that P1 assigns. The random-walk form $T(C) = \sum_{a \in C} v(a)$, with $v(a) \in V_7$ unit vectors, models that directional aggregation and is the tractable model used in the large-deviation analysis of Computation 75. Neither is the unique form of T; an admissible T pairs a magnitude satisfying (14)–(15) with a v-inherited V_7-direction, and its precise form is contingent rather than derived from P1–P3 (§14.3).

The underdetermination this leaves is precise, and bounded. Constrained only by (14)–(15), the magnitude ranges over a family whose dimension grows as $2^{|D|}$. Any order-preserving function on the configuration lattice that vanishes exactly on the distinction-free configurations is admissible, with T_{ex} one representative among them, so the form is not a finite-parameter object (Computation 126). What the constraints fix is not a form but three structural facts: the codomain, T being V_7-valued as forced by the Hurwitz constraint in P1; the magnitude's zero locus (14); and the existence of a nonzero global bias $T(D)$, hence a well-defined direction $\hat{\tau}$ (16). Almost every later derivation reads T through only two of these: the *ordering* of the magnitudes $|T(C)|$, which is what the threshold dynamics of the following chapters require, and the single global direction $\hat{\tau} = T(D)/|T(D)|$; no result uses the direction of $T(C)$ at any non-maximal configuration. *One derivation is not of that kind and must be named*: the Born-consistency comparison (98) tests $\tau_0(C)$ against

Born weights carrying their own normalisation, which is a cardinal read and not an ordinal one (§10.3.2.5). The remaining freedom is therefore confined to exactly the per-configuration content that §14.3 reads as contingent input, which leaves the structural skeleton the rest of this work relies on fully fixed. Computation 140 verifies that robustness by exhibiting the exact invariance under which the instantiation sequence, the per-configuration phase, and the bias direction are all unchanged: monotone reparametrisation of the magnitudes with $\hat{\tau}$ fixed. The one cardinal read holds for the realised content as well, since P3's typicality reaches that content (§1.7, Computations 193, 195), so the conclusion holds in full: only the ordering of the magnitudes and $\hat{\tau}$ are load-bearing, and the canonical form of $T(C)$ is a bounded underdetermination that is no longer carried on the structural front (Appendix C).

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.

4.2. The boundary condition on $|T(C)|$

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

$$|T(C)| \geq \tau \implies C \text{ cannot remain uninstantiated} \tag{18}$$

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

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

Third, the value τ is not a free parameter of the theory. It cannot be chosen arbitrarily or tuned to match observations. As the variational treatment below makes clear, τ is the unique value at which the modal structure of the substrate

changes its minimal configuration from uninstantiated to instantiated. It is derived from the structure of the functional, not imposed from outside.

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

$$\varepsilon(C) \;=\; |T(C)| - \tau \tag{19}$$

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 Chapter 10, 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 book. 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.

4.3. 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 $\psi \to -\psi$ 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.

4.4. 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 § 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 $\psi \to -\psi$ 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 $|\nabla_{\mathcal{C}} \psi|^2_{V_7}$.

Taylor expansion near the threshold. [186] 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 towards $\psi \neq 0$. 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 $|\psi(C) - \psi(C')|_{V_7} \to 0$ as $|C \triangle C'| \to 0$, 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 $|\nabla_{\mathcal{C}}\psi|^2_{V_7} = \frac{1}{4}\sum_{C\sim C'}\langle\psi(C) - \psi(C'), \psi(C) - \psi(C')\rangle$ (the normalisation fixed in (21)) 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 [60] 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 [182]; the differential structure is a further architectural commitment described below.

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

To be precise about what is assumed: in order to write $|\nabla_{\mathcal{C}}\psi|^2$, one needs (a) a topology on $\mathcal{C}$ that makes it possible to speak of nearby configurations, (b) a differentiable structure on $\mathcal{C}$ that makes directional derivatives well-defined, and (c) a metric (or at least an inner product) on the tangent spaces of $\mathcal{C}$ that makes the squared norm meaningful. None of these three items is given for free by the pair (D, δ). The property domain D is a set; its power set

$\mathcal{C} = \mathcal{P}(D)$ inherits only the lattice structure of subsets, not a differential or metric structure. The gradient term therefore represents a genuine extension of the foundational ontology.

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

What the symmetric-difference metric provides is, in fact, all three requirements, (a), (b), and (c), once the Bernoulli measure μ (derived in Chapter 2, §2.5) 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

$$(\partial_a \psi)(C) \;:=\; \psi(C \cup \{a\}) - \psi(C \setminus \{a\}), \tag{20}$$

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 $(\mathcal{C}, d_\triangle)$: two configurations are adjacent in $(\mathcal{C}, d_\triangle)$ 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:

$$|\nabla_{\mathcal{C}} \psi|^2(C) \;:=\; \tfrac{1}{4} \sum_{a \in D} |\psi(C \cup \{a\}) - \psi(C \setminus \{a\})|^2. \tag{21}$$

L^2 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 $\mathbb{E}_\mu[|\nabla_{\mathcal{C}}\psi|^2]$ is finite, because each $\partial_a\psi$ is bounded in L^2 whenever ψ is. The gradient operator $\nabla_{\mathcal{C}} : L^2(\mathcal{P}(D), \mu) \to \bigoplus_{a\in D} L^2(\mathcal{P}(D), \mu)$ is a bounded linear operator, and the Sobolev space [159] $W^{1,2}(\mathcal{P}(D), \mu) = \{\psi : \mathbb{E}_\mu[|\nabla_{\mathcal{C}}\psi|^2] < \infty\}$ is a Hilbert space with inner product $\langle\psi, \phi\rangle_{1,2} = \langle\psi, \phi\rangle_{L^2} + \sum_a \langle\partial_a\psi, \partial_a\phi\rangle_{L^2}$ [56]. 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 $d_\triangle$.

Poincaré inequality. The Efron–Stein inequality [55] states: for any $\psi \in W^{1,2}(\mathcal{P}(D), \mu)$,

$$\mathrm{Var}_\mu(\psi) \;\le\; \tfrac{1}{4}\sum_{a\in D} \mathbb{E}_\mu\big[|\partial_a\psi|^2\big] \;=\; \mathbb{E}_\mu\big[|\nabla_{\mathcal{C}}\psi|^2\big]\,, \tag{22}$$

with equality when $\mu = \bigotimes_a \mathrm{Bern}(1/2)$ and ψ is linear in the indicator functions $\mathbf{1}_a$ [56]. This Poincaré inequality guarantees that the gradient term in $\mathcal{F}$ controls the variance of ψ about its mean: configurations far apart in $d_\triangle$ are forced to have similar ψ values by the cost $c\,|\nabla_{\mathcal{C}}\psi|^2$, 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 Chapter 2, §2.5. The inner product on tangent directions (c) is the $L^2(\mu)$-inner product on Boolean derivatives, derived from $d_\triangle$ 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. Had no derivation of the differential structure from (D, δ) been available, the gradient term would have had to be motivated by a different argument or removed, and the connection to the Einstein–Hilbert action [1, 157] (which relies on the gradient term projecting to the Ricci scalar [10] in Chapter 7) would have needed re-examination. The resolution above discharges this commitment: the differential structure is supplied by $d_\triangle$ together with the $L^2(\mu)$ Hilbert-space structure, both derived from the primitive.

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 exactly what the resolution above supplies; what the gradient term adds beyond it is the coherence requirement itself.

Assembling these three terms gives the modal potential functional. This is the working form of postulate P3 stated in Chapter 1 (equation (1), $\mathcal{F}[\psi, \varepsilon] = \int[-\varepsilon\langle\psi,\psi\rangle + \frac{1}{4}\langle\psi,\psi\rangle^2 + c|\nabla_{\mathcal{C}}\psi|^2]\, d\mu$): under the radial reduction $\langle\psi,\psi\rangle \to \psi^2$, the V_7-valued postulate collapses to the scalar form below. PST adopts the standard Landau–Ginzburg form (quadratic plus quartic) as its postulate; the quartic coefficient $\frac{1}{4}$ is the coefficient of the doublet bilinear $\langle\psi,\psi\rangle^2$ in the doublet convention (§1.7), identified at the matched scale with the Standard Model (SM) doublet form $V_{\mathrm{SM}} = \lambda_{\mathrm{SM}}(\Phi^\dagger\Phi)^2$ via $\langle\psi,\psi\rangle \leftrightarrow \Phi^\dagger\Phi$, and is preserved by the radial reduction.

$$\mathcal{F}(\psi, C) = \int_{\mathcal{C}} \left[a(T(C))\,\psi^2 + b\,\psi^4 + c\,|\nabla_{\mathcal{C}}\psi|^2\right] d\mu(C) \tag{23}$$

where $\nabla_{\mathcal{C}}$ denotes the gradient over configuration space (not spacetime), μ is the canonical Bernoulli measure of equation (9), and $|\nabla_{\mathcal{C}}\psi|^2$ is the

Boolean gradient already defined at (21): the *bare* sum $\frac{1}{4}\sum_{a\in D}|\partial_a\psi|^2$ over the elementary directions, which is also the form (22)'s Efron–Stein equality requires.

That definition has a consequence for c, and it is the opposite of the natural guess (Computation 177). Because the sum is bare while μ is a probability measure, the gradient term is extensive in $|D|$ where the quadratic and quartic terms are intensive, and composing it with (57) gives an emergent kinetic coefficient $K = c\,|D|\,d_0^2/16$. A1 sends $|D| \to V_M/a_n^4 \to \infty$ with d_0 fixed, so a constant c would make the projected field infinitely stiff. *So c is not a constant.* It carries a $1/|D|$, equivalently $c \sim a_n^4 \to 0$, and the $|D|$-independent quantity P3 owns is $\tilde{c} = c\,|D|$, in terms of which $K = \tilde{c}\,d_0^2/16$ is finite. Nothing in P3's content changes by saying so; what changes is that the coefficient is read as a scaling rather than as a number, and that no bare value can be asserted of c at all. The three terms then carry the following interpretation. The quadratic coefficient is $a(T(C)) = -\varepsilon(C) = \tau - |T(C)|$, so that the sign of a tracks the threshold crossing directly. Restricting to spatially uniform configurations ($\nabla_{\mathcal{C}}\psi = 0$) with $b = \frac{1}{4}$, the functional reduces to the compact scalar form:

$$\mathcal{F}(\psi) = -\varepsilon\,\psi^2 + \frac{1}{4}\,\psi^4, \qquad \varepsilon = |T(C)| - \tau \tag{24}$$

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 $\psi^* = \pm\sqrt{2\varepsilon}$. The full three-term structure of equation (23) is retained throughout the analysis.

The first term, $-\varepsilon(T(C))\,\psi^2$, governs the stability of the uninstantiated phase. The coefficient $-\varepsilon = \tau - |T(C)|$ changes sign at $|T(C)| = \tau$. When $|T(C)| < \tau$, $\varepsilon < 0$ 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 $|T(C)| > \tau$, $\varepsilon > 0$ and the term now decreases $\mathcal{F}$ near $\psi = 0$, destabilising the uninstantiated state and driving the configuration towards nonzero ψ. The sign change in ε is precisely what defines the threshold.

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

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

4.5. 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 towards 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 [170] 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, and its sign change at the threshold 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 $c|\nabla_{\mathcal{C}}\psi|^2_{V_7}$ penalises rapid variation of ψ across configuration space. Distributing $\delta/\delta\psi$ across these three terms and collecting, the modal stationarity equation (the Euler–Lagrange equation of $\mathcal{F}$; not an equation of state in the thermodynamic sense of Chapter 16) is the V_7-vector equation:

$$-2\varepsilon\,\psi + 4b\,\langle\psi,\psi\rangle\,\psi - 2c\,\nabla^2_{\mathcal{C}}\psi = 0 \tag{25}$$

This equation has two qualitatively distinct regimes, separated by the thresh-

old.

Below the threshold, $|T(C)| < \tau$: $\varepsilon < 0$. The only solution to the uniform equation is $\psi = 0 \in V_7$. Since $\delta^2\mathcal{F}/\delta\psi^2\big|_{\psi=0} = -2\varepsilon \mathbb{1}_{V_7} > 0$ (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 stationarity equation 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

$$\begin{aligned} \mathcal{V} \;&=\; \big\{\, \psi \in V_7 \;:\; |\psi| = r_0 \,\big\} \;\cong\; S^6, \\ r_0 &= \sqrt{\frac{|T(C)| - \tau}{2b}} = \sqrt{\frac{\varepsilon(C)}{2b}}. \end{aligned} \tag{26}$$

Modal sublimation locks $\hat{\psi}$ to the substrate-determined direction $\hat{\tau} = T(D)/|T(D)|$ on $\mathcal{V}$ (P2). The two "symmetric minima at $\psi = \pm r_0$" are the antipodal pair $\pm r_0\hat{\tau}$ in the 1-dim projection along $\hat{\tau}$; the full vacuum manifold has six tangent directions on S^6 orthogonal to $\hat{\tau}$, five of which, those perpendicular to the complex slice, carry the gauge-theoretic content (§14.2). The appearance of $\varepsilon(C)$ under the radical is structurally significant. The depth of instantiation, how far from zero the order parameter settles, is not a free parameter but is determined directly by how far the configuration's tension sits above the threshold. A configuration with ε barely positive instantiates shallowly; one deep in the instantiated regime carries a large order parameter. This relationship between excess tension and the degree of instantiation is one of the central quantitative predictions of PST.

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

4.5.1. Structural versus spontaneous asymmetry

PST has spontaneous symmetry breaking in the strict sense at the discrete Z_2 level: the $\psi \to -\psi$ symmetry of the functional is broken by the configuration's choice of one of two equally-deep minima. In the richer complex-scalar formulation (and *a fortiori* the full V_7 formulation), the Landau–Ginzburg (LG) potential also carries a continuous symmetry. Its degenerate minima form a U(1) ring in the complex slice, and a G_2-orbit 6-sphere $\mathcal{V} \cong S^6 \subset V_7$ in the full setting. In standard physics that would be the canonical spontaneous symmetry breaking (SSB) setup, producing Goldstone modes along $\mathcal{V}$.

PST departs from the standard SSB account at this angular level. P2's asymmetric vector tension supplies a structural preferred direction $\hat{\tau} = T(D)/|T(D)| \neq 0$ in V_7 *before any vacuum exists*, and modal sublimation *locks* $\hat{\psi}$ to $\hat{\tau}$ by direct determination from the substrate's directional bias (§1.7, P3 statement). The angular direction of the vacuum is therefore *not* a spontaneous choice: it is pinned by P2, upstream of the LG dynamics. What standard physics treats as a historical contingency ("which vacuum did the universe end up in?") PST answers structurally: the substrate's pre-existing tension direction selects it.

Two complementary readings of this layering, both correct:

1. PST inherits the SSB *mathematics*, sombrero potential, degenerate vacuum manifold, Goldstone modes, Higgs radial mode, the electroweak (EW) eaten

Goldstones at the doublet level (§13.4), in full. The discrete Z_2 sign choice in the present paragraph is genuinely spontaneous.

2. PST *redirects* the spontaneous arbitrariness of the angular (continuous) direction: that selection is determined by P2, not by the SSB process. In standard-physics vocabulary the substrate's structural asymmetry is closer to an *explicit* symmetry breaking term, except that the symmetry was never present in the substrate's definition to begin with; P2 is not a perturbation added to a symmetric structure but a postulate of the substrate's structure itself.

The asymmetry is therefore not derived away in PST: it is relocated from the spontaneous-arbitrary level (where standard physics leaves it) to the structural-postulate level (P2). Whether this is progress depends on whether one finds the postulate that the substrate's tension functional is asymmetric under complementation more natural than a random vacuum selection at the SSB step.

One residue of spontaneity survives this relocation, and it is worth being exact about what it costs. The Z_2 choice between $+r_0\hat{\tau}$ and $-r_0\hat{\tau}$ is spontaneous in the standard sense: nothing upstream selects it. It is not, however, an openness of the same kind as the angular selection just removed. The modal potential is invariant under the full G_2 symmetry of V_7, and the involution $\psi \to -\psi$ is a G_2-equivariant special case of that invariance, as noted above. The two minima are therefore related by a symmetry of the entire structure in which they sit, so no quantity definable within that structure distinguishes them. There is no fact as to which sign was taken, and correspondingly nothing for the theory to explain: the residue is not an unanswered question but a question with no content.

The distinction is between two kinds of tie. A tie over whether any distinction obtains at all would be fatal, and it is precisely what the substrate's asymmetry excludes: that is the balanced case, in which nothing is instantiated. A tie over which label an already-drawn distinction carries is harmless, because

the distinction has been drawn and only the naming is undetermined. PST's founding commitment rules out the first. It never undertook to rule out the second, and it does not need to.

The sign of a is 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.

4.6. 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. Restricting again to spatially uniform configurations ($\nabla_{\mathcal{C}}\psi = 0$, the gradient term contributing only a non-negative Laplacian eigenvalue that does not affect the sign at $\psi = 0$) and 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:

$$\left.\frac{\delta^2\mathcal{F}}{\delta\psi^2}\right|_{\psi=0} = -2\varepsilon(|T|)\,\mathbb{1}_{V_7} = 2\,(\tau - |T|)\,\mathbb{1}_{V_7} \tag{27}$$

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 $\delta^2\mathcal{F}/\delta\psi^2\big|_{\psi=0} = 0$, which from the expression above gives $-2\varepsilon = 0$ and thus $|T| = \tau$, a circularity that conceals rather than reveals. The correct approach is to observe that τ is defined by a property of the functional, not by an algebraic equation: it is the 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:

$$\tau = \inf\left\{|T(C)| : \left.\frac{\delta^2\mathcal{F}}{\delta\psi^2}\right|_{\psi=0} = 0\right\} \tag{28}$$

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, its form 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.

4.6.1. Notation: which threshold τ names

An unsubscripted τ means *this* threshold throughout the book: the modal threshold, a condition on the global tension $T(C)$ of a whole configuration. A decoration always signals something else. $\hat{\tau} \in S^6$ is the threshold *direction* that sublimation locks, a unit vector rather than a value. The null threshold τ_{null} (§4.10) is a condition on the *pairwise* tension $T(\{a, b\})$, and it marks the causal boundary: the value at which a separation is null. It is the pairwise companion of τ, determined jointly with it rather than posited independently. The outcome-dependent thresholds τ_i (§10.3) are the modal threshold shifted by a measurement coupling, one per available outcome. Where the text says "the threshold" with no qualifier, the modal threshold is meant.

The passage from $\psi = 0$ to the bifurcated minimum is modal sublimation, with no temporal unfolding. Figure 1 shows the modal potential on either side of the threshold and at it: a single valley below, the flattened critical point at $\varepsilon = 0$, and the two degenerate minima above. §4.8 characterises that transition in detail.

4.7. Summary of the threshold half

Not every tension imbalance produces a universe. The threshold half of this chapter 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

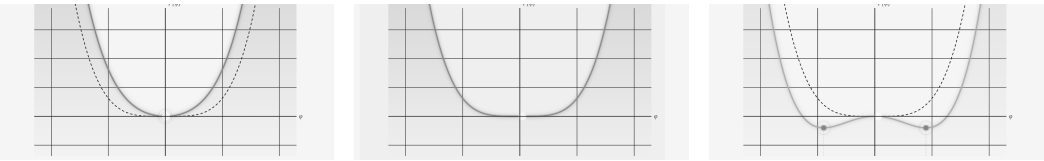


Figure 1: The modal potential $F[\varphi]$ at three values of excess tension ε. *Left:* $\varepsilon = -0.8$, single valley at $\varphi = 0$; the substrate configuration rests at the unique minimum, uninstantiated. *Centre:* $\varepsilon = 0$, the modal threshold; the valley flattens and the minimum is about to split. *Right:* $\varepsilon = +0.8$, two degenerate minima at $\varphi = \pm\sqrt{\varepsilon}$ (the plotted curves use the rescaled variant $-\frac{1}{2}\varepsilon\,\varphi^2 + \frac{1}{4}\varphi^4$; the canonical form of relation (1) below has minima at $\pm\sqrt{2\varepsilon}$); 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 $\Phi : \mathbf{S} \to \mathbf{G}$ from the precausal substrate to instantiated geometry.

the mathematical potential changes from single-welled to double-welled, and its existence is a structural necessity rather than a postulate.

The two postulates in play across this half are P2 (asymmetric vector tension, with $T(D) \neq 0$ in V_7 generically) and P3 (the G_2-invariant Landau–Ginzburg modal potential on V_7); from them follow the existence and value of the modal threshold τ as the bifurcation point of $F[\psi, \varepsilon]$ and the vacuum manifold $\mathcal{V} \cong S^6 \subset V_7$ for $\varepsilon > 0$. One point is worth disclosing explicitly: at the angular level the vacuum direction is *not* a spontaneous choice but is pinned by P2's $\hat{\tau}$, while at the discrete Z_2 level the symmetry breaking is genuinely spontaneous in the standard sense; these layered roles are made explicit in the "structural versus spontaneous asymmetry" paragraph. The modelling choice is the G_2-invariant quartic functional form. The Bernoulli measure, under which F is varied, is not a further choice: it is derived in §2.5 from per-property complementation invariance.

4.8. The character of the threshold

What obtains at 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 it 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 process could unfold, and no space across which it could propagate. What holds at the threshold is not a process but a modal reorganisation: a supra-threshold configuration *is* instantiated as geometry, where a sub-threshold one is not, and the two descriptions stand in this relation tenselessly.

Modal sublimation names this tenseless instantiation relation between the substrate description and the geometric description; it is not an event connecting two times. The substrate's priority over geometry is logical and structural, not temporal: a supra-threshold configuration cannot remain uninstantiated by logical necessity (§17.8), and that logical priority, not a temporal one, is exactly what "precausal" means. Apparent processual language throughout this book, "crossing the threshold," "the transition," "the order parameter relaxes," "the system settles", is shorthand for that relation; no temporal process at the substrate level is implied.

The name *sublimation* is 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 is instantiated *with* time, not located *within* a pre-existing time

- **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.

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, while instantiation is what constitutes that domain, not an event occurring within a prior time. 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 is not located *within* time; time is carried as a consequence of instantiation, not presupposed as a container in which a transition could unfold.

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 Chapter 5.

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 [21], 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 Chapter 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.

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 two modes stand in a strictly modal relation, not a temporal one, even though the phrasings "before" and "after" used here strictly apply only after time has come into being. This is the precise sense in which PST's fundamental relation is irreducible to any physical process: physical processes are changes of state within the instantiated mode, while modal sublimation is the relation that constitutes the mode itself, not a happening within it.

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. Figure 2 traces the whole passage across eight values of the excess tension ε, from the precausal bowl to the degenerate vacuum manifold.

4.9. Categorical formulation: the sublimation functor

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 $f : C \to C'$. The **geometric category G** has as objects Lorentzian manifolds (M, g) and as morphisms smooth maps preserving the causal structure.

Modal sublimation is the functor:

$$\Phi : \mathbf{S} \to \mathbf{G} \tag{29}$$

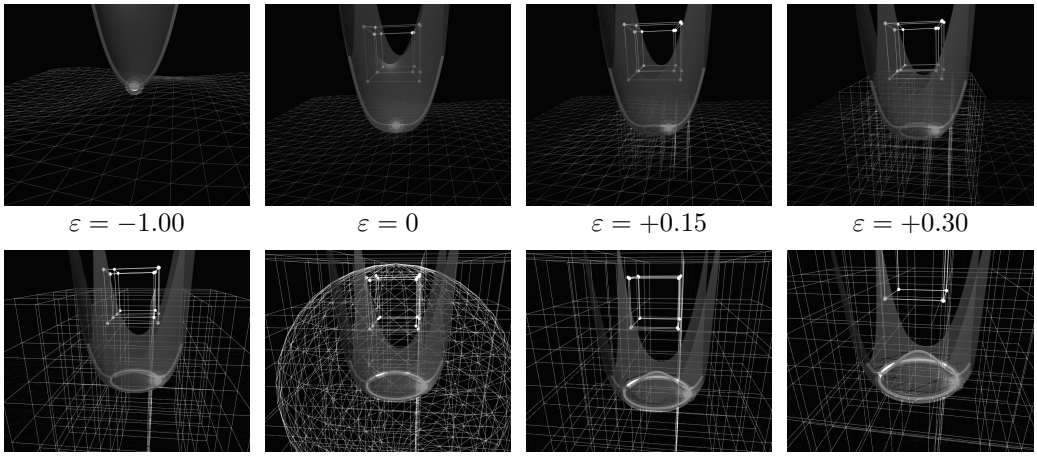


Figure 2: Visual model of modal sublimation across eight values of the excess tension $\varepsilon = |T(C)| - \tau$ (§4.6): $\varepsilon < 0$ precausal, $\varepsilon \geq 0$ instantiated. Sampling is dense on the instantiated side, where the geometric content emerges; a single precausal reference frame 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 bright 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 towards 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 a bright 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.

with four properties encoding the fundamental features of the transition:

1. Φ exists at the categorical boundary; it is not a morphism within **S** or **G**, formalising the nonmechanistic character of the transition
2. Φ carries tension into geometry: the pairwise tensions fix the metric g of the instantiated manifold via equation (30), and the tension projection Π (equation (32)) supplies the stress-energy $T_{\mu\nu}^{(\mathrm{iso})}$; scalar curvature then follows from g through the field equations of Chapter 7
3. Φ is surjective on objects: every Lorentzian manifold of the admissible class, four-dimensional with the topology the spectral-triple split fixes (§1.5), arises as the image of some precausal configuration
4. Φ has no inverse: once a configuration is instantiated, the geometric domain has no access to the precausal substrate that produced it

4.10. 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 realisation. The manifold M and the map ρ are not given independently before the construction; they are jointly determined by it.

A legitimate concern requires attention before the construction proceeds. The notation $\rho : D \to M$ appears to presuppose M, a set of spacetime points, before the construction that is supposed to *produce* spacetime has been carried out. If spacetime points are already available as the target of ρ, then the

theory has not derived spacetime from the substrate; it has merely labelled pre-existing points with properties from D. This would be a foundational circularity.

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

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

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

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

Step 2, Tension-to-metric map. Introduce the null threshold $\tau_{\text{null}} \in \mathbb{R}^+$ as the pairwise-tension companion of τ, determined jointly with τ rather than posited as an independent input. 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. Their relationship reduces to the same input as the rate-integral computation of §16.7: an explicit form of the global tension functional $T(C)$ as a functional of the pairwise tensions $T(\{a, b\})$ at the postulate level, which the canonical form fixes up to the configuration's content (§19.9.1). Once the explicit functional is specified, τ is the critical value of $T(C)$ at which the modal potential loses its single-well stability (§4.1) and τ_{null} is its pairwise restriction; the two are then determined jointly, not independently. The metric is then:

$$g_{\mu\nu}(\rho(a), \rho(b)) = \eta_{\mu\nu} \cdot \operatorname{sgn}(\tau_{\text{null}} - T(\{a, b\})) \cdot |T(\{a, b\}) - \tau_{\text{null}}|^{1/2} \qquad (30)$$

where $\eta_{\mu\nu}$ is a placeholder for the diagonal sign pattern $\operatorname{diag}(+1, -1, -1, -1)$ 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, downstream of the tension structure, by hyperbolic well-posedness of the Goldstone wave equation, and $\eta_{\mu\nu}$ names that pattern: the mostly-minus signature $(+, -, -, -)$ used throughout. The Clifford convention

is $\gamma_0^2 = +I$, $\gamma_i^2 = -I$, giving the local Dirac algebra $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ (Chapter 13). 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, §4.12): the Cauchy problem for $\Box_g \theta = 0$ 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 (30) 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 (30) along the chord, analogous to recovering a Riemannian metric as the Hessian of a distance function. Second, the magnitude $\left|T(\{a,b\}) - \tau_{\text{null}}\right|^{1/2}$ is the minimal-regularity choice. A linear dependence would give a metric with a non-smooth gradient at the null cone, while the square root yields Hölder-$\frac{1}{2}$ behaviour there. That is the appropriate regularity class for a structure that degenerates at $T = \tau_{\text{null}}$, and it ensures that equation (109) blows up precisely at the modal null threshold rather than at $T = 0$.

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

The projection operator. Define the tension projection operator Π:

$$\Pi(T(C))(x) = \int_{\mathcal{C}} T(C') \cdot K(x, C')\, d\mu(C') \tag{31}$$

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 bridges an underlying non-spatial structure to the instantiated domain (the metric itself is supplied by Φ/ρ, not by Π). The stress-energy tensor at $x \in M$ in the isotropic, homogeneous case is:

$$T_{\mu\nu}^{(\mathrm{iso})}(x) = \Pi(T(C))(x) \cdot g_{\mu\nu}(x) \tag{32}$$

This scalar form $T_{\mu\nu}^{(\mathrm{iso})} = \Pi(T(C))\, g_{\mu\nu}$ 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. Two facts fix its scope exactly, and they are narrower than "isotropic" (Computation 155). Applying the Bianchi identity to $T_{\mu\nu} = \Pi\, g_{\mu\nu}$ gives $\nabla^\mu T_{\mu\nu} = \partial_\nu \Pi$, so conservation *forces* Π to be constant; and a perfect fluid is proportional to $g_{\mu\nu}$ if and only if $\rho + p = 0$, i.e. $w = -1$. So the scalar form is not the leading-order approximation for isotropic sources: an isotropic dust or radiation fluid is already outside it. It is exactly and only the cosmological constant. 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. §7.4 supplies its structure, the mechanism that sets the equation of state, and the weight itself, which is fixed by requiring the tensor and scalar projections to agree (§7.4.4). 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.

4.11. Uniqueness of Φ and the origin of diffeomorphism invariance

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

The proof follows directly from the construction. Given ρ, the metric $g_{\mu\nu}$ at every pair of points is fully determined by equation (30): there are no further choices once the realisation map is fixed. Now suppose $\rho' = f \circ \rho$ for some smooth bijection $f : M \to M'$. The tension values $T(\{a, b\})$ depend only on $a, b \in D$, not on their images under ρ, so they are unchanged by the substitution. Therefore the metric computed at $(\rho'(a), \rho'(b))$ via equation (30) 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.

What the uniqueness theorem leaves open is not a further gauge choice but *which* substrate configuration is instantiated: the realisation map ρ is fixed up to diffeomorphism once the configuration is given, but the theory must still say which configuration sublimation realises. P3 supplies this through maximum entropy: modal sublimation instantiates the μ-typical (maximum-entropy) configuration, not a μ-exceptional one. The permutation symmetry of the substrate measure already forces equidistribution *in measure*; maximum-entropy sublimation strengthens this to the realised configuration itself, so its emergence cloud ρ_n is metric-covariantly equidistributed at the closure

rate $\eta_{K,R} < d_0^2/4$. This is exactly condition A3 of the Mosco limit (§18.1); it therefore holds as a consequence of P3, and the Mosco convergence to $M = \mathbb{R} \times S^3$ closes under A1–A7, A3 being supplied by P3 rather than assumed (§18.1). The selection is a typicality principle internal to P3, the analogue of maximum-entropy or generic-initial-data selection elsewhere in physics: it adds no fourth postulate and derives no rate from more primitive structure, but it locates the equidistribution where it belongs, in the dynamics of instantiation.

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, which is the uniqueness theorem stated at the head of this section.

A note on the metric signature. The argument in Step 2 assigned the Minkowski signature $(+,-,-,-)$ through equation (30), 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, whose standard form 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. 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 beyond the Gram construction. The Lorentzian signature $(+,-,-,-)$ is fixed as a *theorem* via Computation 2 of §4.12: 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 real K-theory (KO) additivity given the verified spectral-triple structure (§1.5): 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, δ) + LG + the realisation map ρ, with both signature and macroscopic dimension outputs rather than inputs.

4.12. Lorentzian signature from Goldstone hyperbolic well-posedness

This section makes self-contained the result that Lorentzian signature follows from hyperbolic well-posedness of the Goldstone wave equation (Computation 2): we state the partial differential equation (PDE), the function spaces, the energy estimate, and show why one timelike direction (signature $(p, q) = (1, n-1)$) is the unique signature compatible with a well-posed Cauchy problem on the projected spacetime $\mathbb{R} \times S^3$.

The logical order forestalls any apparent circularity. The manifold $M = \mathbb{R} \times S^3$ and its differentiable structure emerge first, as the Mosco limit of the Boolean Dirichlet form (§7.8), carrying a topology but as yet no metric signature; the signature is then fixed downstream as the unique one for which the Goldstone wave equation on that already-emergent manifold admits a well-posed Cauchy problem. Hyperbolicity does not presuppose a metric, it selects among the candidate signatures available on a manifold that exists prior to any of them.

The PDE. After modal sublimation has selected the vacuum manifold $\mathcal{V} \cong S^6 \subset V_7$ (§1.7, P3), the Goldstone phase θ along the broken direction satisfies, at the quadratic order of the LG functional, the second-order PDE

$$g^{\mu\nu}\, \partial_\mu \partial_\nu\, \theta(x) \;=\; 0, \qquad x \in M, \;\; \mu, \nu = 0, 1, \ldots, n-1, \tag{33}$$

where $g^{\mu\nu}$ is the (Wick-rotation-undetermined) [161] signature-(p, q) metric on $M = \mathbb{R} \times S^3$ obtained as the projection of the LG kinetic term to emergent spacetime. The order parameter contributes $r_0^2|\nabla\theta|^2$ to the LG free energy, which under $\nabla \to g^{\mu\nu}\partial_\mu \otimes \partial_\nu$ becomes the wave operator.

A note on which degree of freedom this is. (33) is written massless, as the Goldstone mode of a spontaneously broken symmetry is. The order parameter's *angular* directions are not that: P2 supplies $\hat{\tau}$ before any vacuum exists and modal sublimation determines $\hat{\psi}$ from it (§4), so those six directions are constrained rather than dynamical, and the residual stabiliser $\mathfrak{su}(3) = \mathrm{Lie}\,\mathrm{Stab}_{G_2}(\hat{\tau})$ fixes the vacuum and is therefore unbroken (the orbit is the same S^6 whether the symmetry is read as G_2 or as the SO(7) of the bare sombrero that Computation 170 counts in, so the six constrained directions are the same). Neither yields a Goldstone (Computation 170). *What carries θ is the phase*: $\hat{\tau}$ is a *real* direction in V_7 and determination fixes a direction while saying nothing about a phase, so the complex line $\mathbb{C}\hat{\tau}$ survives the constraint where the real line does not, and the phase of the complexified order parameter is left free (Computation 171). It is exactly massless whenever the couplings to $\hat{\tau}$ are phase-blind, which gauge invariance enforces, since $\mathrm{Re}\langle\psi, \hat{\tau}\rangle$ is not invariant under $\psi \mapsto \mathrm{e}^{i\alpha}\psi$ while the moduli are. That makes the complex-valued order parameter a *requirement* rather than the richer of two formulations: a real one has no phase and so no θ, and with no θ the signature theorem loses the mode it argues from. *What does not depend on the answer* is everything below: Hadamard well-posedness is decided by the *principal* symbol $g^{\mu\nu}\xi_\mu\xi_\nu$, so even a massive θ would leave the argument untouched. The signature theorem needs hyperbolicity, not masslessness, and is more robust than its original statement claimed (Computation 170).

Cauchy problem and function spaces. The Cauchy problem for equation (33) on the time slice $\Sigma_0 = \{0\} \times S^3$ asks: given initial data $(\theta_0, \theta_1) \in H^1(S^3) \times L^2(S^3)$, does there exist a unique solution $\theta \in C(\mathbb{R}; H^1(S^3)) \cap C^1(\mathbb{R}; L^2(S^3))$ with $\theta|_{\Sigma_0} = \theta_0$, $\partial_t\theta|_{\Sigma_0} = \theta_1$?

For the Cauchy problem to be *well-posed in the Hadamard sense* [110] (exis-

tence, uniqueness, and continuous dependence on initial data) the principal symbol $g^{\mu\nu}\xi_\mu\xi_\nu$ of equation (33) must be *strictly hyperbolic* with respect to the splitting Σ_0, i.e. the polynomial

$$g^{00}\omega^2 \;+\; 2g^{0i}\omega\xi_i \;+\; g^{ij}\xi_i\xi_j \;=\; 0 \tag{34}$$

must have exactly two real roots $\omega = \omega_\pm(\xi)$ for every nonzero spatial covector $\xi \in T^*\Sigma_0$ [99]. Strict hyperbolicity forces the metric to have one timelike direction (transverse to Σ_0) and $n-1$ spacelike directions. Equivalently, signature $(p,q) = (1, n-1)$ (or its sign-reversal $(n-1, 1)$, the same content).

Energy estimate. Given strict hyperbolicity, multiplying equation (33) by $\partial_t\theta$ and integrating over Σ_t yields the standard wave energy

$$E(t) \;=\; \tfrac{1}{2}\int_{\Sigma_t} \big(\,(\partial_t\theta)^2 \;+\; g^{ij}\partial_i\theta\,\partial_j\theta\,\big)\sqrt{g_{S^3}}\,dx \tag{35}$$

which is conserved: $E(t) = E(0)$ for all $t \in \mathbb{R}$. Energy conservation immediately gives existence (via the standard semigroup construction [99]), uniqueness (the zero-data Cauchy problem has zero energy and hence zero solution by Gronwall), and continuous dependence (initial-data perturbations propagate at finite speed bounded by $\max|g^{ij}/g^{00}|^{1/2}$ along light cones). This is the standard Sobolev-space well-posedness theory for the wave equation [99].

What fails in other signatures. Three alternatives must be considered and ruled out:

1. *Euclidean* $(p,q) = (0,n)$. All directions are spacelike; the principal symbol $g^{\mu\nu}\xi_\mu\xi_\nu$ is positive-definite. Equation (33) becomes *elliptic* (Laplace) rather than hyperbolic. The Cauchy problem for an elliptic equation is famously *ill-posed* (Hadamard's classical example [110]): small perturbations in the data give exponentially growing perturbations in the solution, so no continuous-dependence map exists. Modal sublimation produces a domain that admits time evolution, not an elliptic boundary-value problem;

Euclidean signature is therefore incompatible with the Cauchy-problem structure of modal sublimation.

2. *Ultra-hyperbolic* (p, q) *with* $p \geq 2$*,* $q \geq 2$*.* Two or more timelike directions. The principal symbol has signature (p, q) as a quadratic form, but the Cauchy problem on the codimension-p "time" submanifold has multiple time derivatives in the data, and strict hyperbolicity in the sense above fails: the polynomial in ω has either no real roots (for some ξ) or more than two. The Cauchy problem is ill-posed in the Hadamard sense; existence fails for generic smooth data [111].

3. $(p, q) = (n, 0)$ *(anti-Lorentzian).* All directions timelike, no spatial direction. Equation (33) has no propagation direction; the projection has no instantiated geometry to speak of. Excluded as a degenerate case.

Conclusion. The unique signature compatible with a well-posed Cauchy problem for the Goldstone wave equation on the Mosco-limit spacetime $M = \mathbb{R} \times S^3$ is $(p, q) = (1, n-1)$. Combined with the $n = 4$ fix from the minimal KO-split (§1.5), this gives $(p, q) = (1, 3)$, i.e. signature $(+, -, -, -)$. Lorentzian signature is therefore a *theorem* of modal sublimation rather than an additional postulate: the directionality of the threshold crossing selects time, hyperbolic well-posedness selects the number of spatial directions, and Computation 2 classifies the admissible signatures symbolically by the principal symbol of the free wave operator, beside its Derrick scaling table.

The above is a structural derivation using standard PDE theory [99]; the energy estimate is rigorous and the alternative-signature arguments are sketched at standard textbook level rather than fully worked. A fully rigorous treatment with detailed Sobolev-space analysis is straightforward but would expand this section substantially; the structural conclusion does not depend on which level of analytic detail one prefers. Computation 2 carries the symbolic signature classification; the energy estimate itself is standard and is not reproduced there.

4.13. 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.

Concretely, the modal-sublimation functor $\Phi : \mathcal{S} \to \mathcal{G}$ is constructed explicitly in three steps. Its central derived result (Computation 2) is that hyperbolic well-posedness of the Goldstone wave equation along $\hat{\tau}$ forces the Lorentzian signature $(+,-,-,-)$. The energy-estimate and PDE details are sketched here, with the energy estimate itself standard and not reproduced in full. The sign cut at τ_{null}, where hyperbolicity flips, exists because that well-posedness requirement demands it, not because it is posited; its value is determined jointly with τ through the explicit form of $T(C)$, the same explicit form the rate integral of §16.7 awaits, so the joint determination is $T(C)$-contingent (C.7) rather than open at the structural level. Diffeomorphism invariance is likewise derived, from the non-uniqueness of the realisation map ρ. The passage $\mathcal{S} \to M = \mathbb{R} \times S^3$ is a theorem under A1–A7, resting on Mosco convergence of the Boolean Dirichlet form to the Laplace–Beltrami form on M. The argument is structural, and is stated explicitly in §18.1 with seven assumptions A1–A7. Equicoercivity A7 is itself a theorem given A1, A2, A4 and A3 of §18.1.6, which combines a standard Poincaré inequality with the equidistribution condition A3, a uniform-rate hypothesis on H^1-bounded compactly-supported test functions, supplied by P3's maximum-entropy sublimation. A3 is corroborated for the i.i.d. surrogate (Computation 70), which confirms the rate P3 supplies for the emergence-determined cloud. The minimal KO-split $4 + 6 = 10$ given the KO-6 substrate (§1.5) fixes $n = 4$, and is

not a selection: KO-dimension is additive and the admissible pair is unique, so there is nothing to minimise over (§1.5, Computation 148).

5. Coinstantiation of Spacetime and Causality

Space is normally pictured as a stage, with cause and effect as things that happen on it. This chapter argues they are not two things at all, but two ways of describing one new structure, and that they appear in the same instant. Neither can come first, because before that instant there was no earlier.

5.1. Causality as coinstantiated structure

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

This point carries weight precisely because it is not obvious. Much of the philosophical and cosmological literature treats causality as prior to or more fundamental than spacetime. In the Kantian tradition, causality is a category of the understanding applied to experience, logically prior to any empirical fact about the world [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 Chapter 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:

$$p \leq q \iff q \in J^+(p, g) \tag{36}$$

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

$$\begin{aligned} J^+(p, g) = \{\, q \in M \;\mid\; & \exists \text{ future-directed causal curve} \\ & \gamma : [0, 1] \to M,\ \gamma(0) = p,\ \gamma(1) = q \,\} \end{aligned} \tag{37}$$

A curve is causal under g if and only if its tangent vector $\dot{\gamma}$ is everywhere non-spacelike, that is $g(\dot{\gamma}, \dot{\gamma}) \geq 0$ 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 (36) has no pre-image in the precausal category **S**. Formally, there is no morphism or relation in **S** that maps to $\leq$ under Φ, because the objects of **S** are configurations equipped only with δ and T. This is not a deficiency of the model; it is a structural theorem: the functor Φ is not surjective on structure but on objects. It produces a manifold but does not carry into that manifold a pre-existing causal order; the causal order is generated anew by the Lorentzian signature of the metric produced by equation (30). Therefore the following holds:

$$\text{causality} \;=\; f(g) \;=\; f(\Phi(C)) \tag{38}$$

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

5.3. 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 precausal substrate. Asymmetric tension, by definition, is a non-symmetric functional: $T(\bar{C}) \neq -T(C)$ generically, where $\bar{C}$ is the complement of C in D. Across the functor Φ this asymmetry survives as the time orientation: the Lorentzian signature itself is fixed by hyperbolic well-posedness of the Goldstone wave equation (Computation 2), and the

preferred time direction of the manifold reflects the substrate's globally fixed modal-sublimation direction $\hat{\tau} \in V_7$, locked at threshold crossing by P3. 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 $T(\bar{C}) \neq -T(C)$ (11) is a configuration-level condition; a global time orientation on M requires a continuous nowhere-vanishing timelike vector field on the entire manifold. A Lorentzian manifold is time-orientable if and only if such a field exists [38], 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 globally fixed modal-sublimation direction $\hat{\tau} \in V_7$ (Chapter 4) gives a precausal structure that Φ maps to a candidate timelike field on M.

The continuity and nowhere-vanishing of this field follow from three structural ingredients.

(i) The modal-sublimation direction $\hat{\tau} \in V_7$ is fixed globally at threshold crossing (§1.7, P3), so it is a single non-vanishing element of V_7 rather than a configuration-dependent choice.

(ii) The substrate-to-spacetime projection ρ underlying Φ is permutation-equivariant under the substrate's S_D symmetry, by the same permutation-invariance argument that drives the cosmological-principle theorem (§16.2, "Theorem (Cosmological Principle)"). That argument also forces the equidistribution assumption A2 of §18.1.

(iii) The Mosco-convergence framework of §18.1 delivers ρ in the strong-resolvent topology, which preserves continuity of the projected direction.

Combining (i)–(iii): $\hat{\tau}$ is mapped by Φ to a globally non-vanishing continuous timelike direction on $M = \mathbb{R} \times S^3$; specifically $\Phi(\hat{\tau}) = \partial/\partial t$ along the $\mathbb{R}$ factor. Geroch time-orientability is therefore a structural consequence of the modal threshold (P3) plus the substrate permutation invariance plus the Mosco

convergence, rather than an additional assumption.

5.4. The causal order and the degenerate vacuum

A further consequence of coinstantiation concerns the vacuum manifold. As established in Chapter 4, 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 book, 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 (§14.2). 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 §9.5.

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 precausal substrate possesses property differentiation, and property differentiation entails asymmetric tension, and once tension exceeds the modal threshold, modal sublimation to instantiated geometry is logically necessary, not contingently triggered. The question Leibniz posed, "why is there something rather than nothing?" [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 tunnelling 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, 232] makes the causal partial order primitive and attempts to derive the continuum from it. It is the sharpest test of this chapter's claim, because a causal set is exactly what §5.1 said does not exist: a causal relation with no spacetime, a locally finite partial order and nothing else. Its warrant is Malament's theorem [223], that in a past and future distinguishing spacetime the causal order alone fixes the topology, the differentiable structure and the conformal metric, so that the order plus a volume count recovers g. 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. The two are not in contradiction, and the reason should be stated rather than assumed. Malament's theorem cuts PST's way: in the continuum the order and the conformal geometry are one structure read twice, which is coinstantiation stated as a theorem. What the discrete orders add is the result of Kleitman and Rothschild [224], that almost every finite partial order is a three-layer object approximated by no manifold at all; the manifold-like orders are a vanishing fraction, and causal set theory needs a dynamics, sequential growth or a sum over orders weighted by a discrete action, to make them typical. In PST's vocabulary the generic finite order is a combinatorial object that is called causal by its name. It carries a causal relation in the sense of this chapter only when it is the order of a Lorentzian metric, which the sprinkling limit supplies and the generic order lacks; the claim of §5.1 is about the relation, not about the name. 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. 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.

Nothing new is postulated in this chapter: causality co-emerges with geometry through the same modal-sublimation step, with no separate causal postulate added, and the arrow of time follows as a structural consequence of P2's asymmetric tension direction $\hat{\tau}$. The causal order itself is interpreted as the partial order on instantiated events induced by Φ, an interpretation consistent with the standard relativistic causal structure on M.

6. The Emergence Chain: Eight Steps to a Universe

This chapter is the route map. Eight steps take the argument from "some things differ" all the way to "there is a universe with matter in it". It is worth reading slowly, because every chapter around it is one of these steps done properly.

6.1. Overview and the eight-step summary table

The preceding four chapters 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 chapter 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.

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.

Step	Concept	Mathematical object
1	Property differentiation	(D, δ, v), a set with a symmetric, irreflexive distinction relation and V_7-valued properties
2	Asymmetric tension	$T : \mathcal{C} \to V_7,\ T(\bar{C}) \neq -T(C)$ generically
3	Modal threshold	Bifurcation point τ of $\mathcal{F}(\psi, C)$, eq. (28)
4	Modal sublimation	Functor $\Phi : \mathbf{S} \to \mathbf{G}$, constructed via ρ and eq. (30)
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	$T_{\mu\nu} = \Lambda\, g_{\mu\nu} + T_{\mu\nu}^{(\text{fields})}$, the vacuum term plus the projected matter fields, with $\Pi_{\mu\nu}$ the conformal part of the second; $G_{\mu\nu} = 8\pi G\, T_{\mu\nu}$; bosonic field sector; fermionic content with structural spin-statistics from the canonical anticommutation relations (CAR) algebra on the Boolean distinction primitive (Chapter 11)
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 $\mathcal{V} \cong S^6$

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 towards 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 and the gradient regularisation) are the minimal motivated choices consistent with the structural requirements. They are not strictly entailed by (D, δ) alone, as the epistemological caveat of §3.3 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 Chapter 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 a configuration and its complement do not carry opposite tension, $T(\bar{C}) \neq -T(C)$ generically, equivalently $T(D) \neq 0$ (P2, (11)). This condition is the precise additional commitment that converts step 1 into step 2. It cannot be derived from (D, δ) alone; it is the assertion that the precausal substrate is generically asymmetric, and it functions as the bridge between the bare existence of distinctions and their structural expression as tension. Granting it, the move from step 1 to step 2 is structurally closed: step 2 is step 1 made structural, with the complement-asymmetry condition as the content of that structuring move. This is the foundational hinge of the chain.

6.4. 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 (28).

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 towards modal coherence (captured by the quadratic restoring term) and the tendency towards 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 $\Phi : \mathbf{S} \to \mathbf{G}$, the map from the precausal category to the geometric category. The functor is not postulated; its existence follows

from the threshold condition together with the definition of the two categories. The precausal category **S** has as objects configurations with tension above τ; the geometric category **G** has as objects Lorentzian manifolds. The functor Φ assigns to each configuration its geometrically realised image, constructed explicitly via the realisation map ρ and the tension-to-metric assignment of equation (30). 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 & 6: sublimation to spacetime & causality

The image of Φ is a Lorentzian manifold (M, g). The Lorentzian signature is not assumed but derived, and not from the threshold cut alone. Equation (30) sorts pairs by their tension about the null threshold τ_{null}: above it to spacelike separations, below it to timelike, at it to null. That fixes the sign of each separation given the cut; it does not fix how many directions fall below the cut. That count, one, is the theorem of hyperbolic well-posedness of the Goldstone wave equation (Computation 2, §4.12). The $(+, -, -, -)$ signature of spacetime is therefore an expression of the way tension distributes across pairwise configurations together with the well-posedness requirement, not of the distribution alone.

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 vacuum term of the stress-energy tensor is the projection of the precausal tension functional through the operator Π, as in equation (39), and the *conformal part* of the matter term is the tensor projection (40), whose weight is derived rather than adopted (§7.4.4). The rest of the matter term is carried by the projected fields rather than by the pair geometry (§7.4.6, Computation 159). What the projection does supply 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. The field equations (48) 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 is derived in Chapter 7 and the bosonic field sector in §12.1, both from the order parameter ψ. The fermionic sector is established by the three results of §11.3: CAR-to-continuum Mosco convergence, Wightman [108] axioms verification, and the substrate-derived Dirac equation as dynamical content. The statistics half of the spin-statistics rule is a structural theorem from the Boolean exclusion of P1,

not a separate input; the spin label is attached by the spinor bundle (§11.3). 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. The Standard-Model Higgs sector closes adjacently on the substrate side. The field-normalisation identity $Z^2 = \mathrm{e}^{-1}$ has its substrate-side factor derived structurally from P1–P3 via tensor-product factorisation forced by P1's Bernoulli product measure (Computation 100, §12.3, §12.5). Identifying that substrate factor with the Standard Model (SM) side ratio $\lambda_{\mathrm{SM}}(M_*)/b$, bridge premise (B), is reduced within PST to a field-strength-renormalisation identity by a wave-function-renormalisation argument and the unique-soft-mode theorem (§12.5). Its internal renormalisation is derived; what the Standard Model (SM) side of that matching still rests on is F4, that the SM is PST's effective field theory (EFT) below M_* (§12.6.2.1), together with the renormalisation-group equation (RGE) test.

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 $\hat{\tau}$. That angle is the phase of the complex line $\mathbb{C}\hat{\tau}$, the one angular degree of freedom the sublimation constraint leaves free (Computation 171); the real directions tangent to S^6 are constrained and carry no Goldstone (Computation 170). 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 and the angular momentum of electrons in atomic orbitals both reduce to the single fact that the residual vacuum in each 2-plane slice 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 sole primitive, P1; P2 and P3 enter at steps 2 and 3 as the other two postulates. The remaining steps follow 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 postulates plus one derived structural statement (§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 book.

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. What PST does not claim to derive is the Standard Model's *contingent* content: the particular masses and couplings, which depend on what C happens to be, in the way that topology does not fix the dimensions of a particular manifold. Its *structural* content is another matter and is derived: the gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ as the unimodular unitaries of A_F, the colour count $N_c = 3$ from the stabiliser of the threshold direction, and one generation's charge assignments (Chapter 13; the full roster is §20.7.1). What PST claims beyond that 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. Summary

This chapter 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 postulates (P1–P3: the distinction primitive, asymmetric tension, and the Landau–Ginzburg modal potential), together with the Laplace-type projection condition, which is derived from P1–P3 under A1–A7 via Mosco convergence (§7.8). 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.

This chapter introduces no new postulates. It draws only on P1, P2 and P3, and each of the eight steps is a structural derivation conditional on the preceding ones. The Laplace-type condition entering Step 5 is derived from P1–P3 under A1–A7 by the same Mosco convergence set out in §7.8. The chapter is integrative rather than a source of new claims; its purpose is to display the logical chain in one place.

7. Energy and Matter as Projected Tension

Einstein's equations say that matter tells space how to curve. What they do not say is where the matter came from; you have to put it in by hand. This chapter argues that in PST it was there all along, and that what is called matter and energy is the substrate's leftover pull seen from inside.

7.1. Results imported from the technical core

Everything from this chapter onwards leans on four analytic results; three are proved in the convergence chapter at the back of the book (§18.1 onwards) and the fourth, (C4), in §8.10. They are collected here, as statements rather than proofs, for one reason: a reader should be able to follow the physics of the chapters that follow without leaving them. **Each result below is used later exactly as stated**, so a later reference to the convergence chapter points at the proof of something already in hand, not at something the reader must go and learn.

(C1) The Mosco limit (§18.1). Under the seven assumptions A1–A7, the Boolean Dirichlet forms on $\mathcal{P}(D_n)$ converge in the Mosco sense to the Laplace–Beltrami form on $M = \mathbb{R} \times S^3$, and the convergence is strong-resolvent, so the Laplace–Beltrami operator is identified as the generator of the limit. The equidistribution assumption A3 is not an extra hypothesis: P3 supplies it as a typicality property of maximum-entropy sublimation. **The limit therefore closes under A1–A7 with A3 discharged by P3**, and what it still rests on beyond P1–P3 is A1, which §18.1 carries as architectural input, together with the convergence-rate control of A6, standard but not yet written out.

(C2) The K-uniform rate (§18.2, with the reduction lemma at §18.1.6). This hypothesis is proved *for the i.i.d. surrogate*, by a five-step Berry–Esseen argument with three closures, and in the flat-bulk case, the curved case following by bounded geometric factors in the Sobolev constants. From that rate, equicoercivity A7 follows by the reduction lemma; A7's own condition A3 is supplied by P3 rather than assumed, so A7 is a theorem given A1, A2, A4

and A3. A1 remains architectural input, so this does not make the chain unconditional on P1–P3 alone. This is what keeps (C1) from resting on a hidden constant. **The surrogate qualification is real and is carried forward wherever the rate hypothesis is used.**

(C3) Dirac-operator convergence (§18.3). The spectral-triple data converge alongside the forms, which is what licenses treating M as a noncommutative-geometric object downstream. **This is the one of the four whose closure is partial**: it is closed at the relaxed-criterion level, and the existence of a finite uniform bound over the operative range $k \leq k_D$ is an analytical theorem, but the sharp constant 3.7 is derived at its maximiser and scanned elsewhere rather than proved over the whole regime. It is now less opaque than that sounds: the Dirac operator drops out of the problem entirely, the case $k = 1$ closes in closed form, the large-N limit of the deficit at fixed k is in closed form, and the sharp value is an explicit algebraic number at the corner $(6, 12)$ (§19.2.1, Computations 157, 213 and 214). Nothing below turns on the sharp value.

(C4) Kernel normalisation (§8.10). The kernel normalisation $\mathcal{N}$ appearing in Newton's constant, $G = 6\pi d_0^2/(c\, r_0^2\, \mathcal{N})$, is the trace $\mathcal{N} = \mathrm{Tr}_{\mathcal{H}_F}(\mathbf{1})$ over the finite internal factor, the fermionic multiplicity of one generation. It is therefore a finite number fixed by the matter content and **not a free parameter**, which is what makes the Einstein–Hilbert coefficient of Chapter 8 a derived quantity rather than a fitted one.

What these four buy, taken together, is the right to say that the projected geometry *is* $M = \mathbb{R} \times S^3$ with its Laplace-type operator, rather than merely resembling it. Every downstream statement that rests on the Mosco limit, the Einstein–Hilbert action from the Seeley–DeWitt expansion, the spectral-triple structure, and the Casimir prediction insofar as it uses the kernel, inherits that standing.

7.2. 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.

7.3. The stress-energy tensor as projected tension

In the isotropic, homogeneous approximation the precise statement is:

$$T_{\mu\nu}^{(\mathrm{iso})}(x) = \Pi(T(C))(x) \cdot g_{\mu\nu}(x) \tag{39}$$

where $\Pi(T(C))(x)$ is the projection of the precausal tension functional onto the point $x \in M$, as defined in equation (31). 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} would reflect the directional distribution of tension across configurations that realise nearby points, and the pressure components T_{ij} the isotropic or anisotropic character of the tension distribution; these directional components lie beyond the scalar weight that Π supplies and require the tensor-valued projection $\Pi_{\mu\nu}$, constructed in the next section. None of these quantities is postulated: the scalar weight is derived from the single precausal input $T(C)$ through the action of Π, and the directional components follow from the same pairwise tensions once the dyad structure is in place.

7.4. The tensor projection $\Pi_{\mu\nu}$

The scalar weight is the trace of something larger, and the larger object is built from material the theory already has. §4.10 fixed the scalar form's scope: it is exactly the cosmological constant, not an approximation for isotropic sources. This section constructs the tensor that contains it, records the one completion that does *not* work, and says what stays open. Everything below is checked in Computation 155.

7.4.1. Why the obvious completion fails

A stress tensor needs a direction per contribution, and PST's tensions are pairwise, so the natural candidate is one dyad per pair: for $\{a, b\}$ with separation $\Delta^\mu = \rho(a)^\mu - \rho(b)^\mu$, take the unit vector $\hat{e}^\mu = \Delta^\mu / \sqrt{|\Delta \cdot \Delta|}$ and the dyad $\hat{e}_\mu \hat{e}_\nu$. That normalisation diverges as $\Delta \cdot \Delta \to 0$, which is precisely where τ_{null} sits, since (30) puts the null condition at $T(\{a, b\}) = \tau_{\text{null}}$. The candidate is singular exactly on the surface the whole construction is built around. Measured: over clouds of 200 to 800 realisation points the per-pair value of the summed dyad swings by 56% and does not settle, because a handful of near-null pairs dominate it.

That is worth recording rather than passing over. The light cone is not an awkward region of this construction; it is the construction's own zero, and any completion that normalises by the interval will fail there.

7.4.2. The separation dyad and its trace

Drop the normalisation, and weight the unnormalised dyad by a function of the pair's own invariant:

$$\Pi_{\mu\nu}(x) \;=\; \tfrac{1}{2} \sum_{a \neq b} W(s_{ab})\; \Delta_\mu\, \Delta_\nu\; B(x;\, \rho(a),\, \rho(b))\,, \qquad s_{ab} = \Delta \cdot \Delta, \tag{40}$$

with B a bond kernel supported between the two realisation points. Nothing is normalised, so nothing is singular. And the trace is immediate: $g^{\mu\nu}\Delta_\mu\Delta_\nu = s_{ab}$, which (30) ties directly to $T(\{a, b\}) - \tau_{\text{null}}$. The scalar Π of (31) is therefore not lost: it reappears as the *trace* of (40) rather than as a factor multiplying the metric. That is the structural correction the previous section's result demands. A scalar times $g_{\mu\nu}$ can only be Λ; a scalar as a trace constrains nothing.

7.4.3. What fixes the equation of state

If W depends only on the invariant, then on an ensemble rotationally isotropic about a rest frame u^μ the sum (40) can only take the form

$$\Pi_{\mu\nu} \;=\; \rho\, u_\mu u_\nu \;+\; p\, h_{\mu\nu}, \qquad h_{\mu\nu} = u_\mu u_\nu - g_{\mu\nu}, \tag{41}$$

which is the perfect fluid. It is not imposed but forced by the ensemble's symmetry, and the residual of the decomposition falls from 1.5×10^{-1} to 4.8×10^{-2} as the cloud grows from 200 to 1600 points.

The two coefficients then have a reading the theory already supplies. ρ is the moment over *timelike* pairs and p the moment over *spacelike* ones, and those two classes are separated by $T(\{a, b\})$ against τ_{null}, the same cut (30) makes to fix the Lorentzian signature. So the threshold that gives the light cone gives the equation of state, and $w = p/\rho$ is the balance between the tension carried on spacelike pairs and on timelike pairs. Nothing else enters.

Writing the weight as $+f(|s|)$ on timelike pairs and $-\lambda f(|s|)$ on spacelike ones makes λ that budget ratio and the only free scalar. Then $w = -1$ is one scalar condition, $\rho + p = 0$. Solving it in the profile family tested gives $\lambda^* = 0.8699$, and there the measured $\Pi_{\mu\nu}$ is proportional to $g_{\mu\nu}$ to 3.7×10^{-2}, which is exactly the decomposition residual, so at λ^* the tensor *is* the scalar form. That is the reading inside a family the theory does not have: the next subsection fixes the weight, the derived weight gives $w = 1/3$ exactly for this sector (Computation 158), and the cosmological constant of §16.5 is the separate vacuum term, never a point of the pair sum (§7.4.6).

7.4.4. The weight is not free either

The budget ratio above was written as though W could be chosen. It cannot. Two objects the theory already has, made to agree, fix it completely (Computation 158).

(31) projects the tension itself, and the trace of (40) is $g^{\mu\nu}\Delta_\mu\Delta_\nu = s$. (30)

ties the two together: $|s| = |T - \tau_{\text{null}}|^{1/2}$ with $\text{sgn}(s) = \text{sgn}(\tau_{\text{null}} - T)$, so

$$T(\{a,b\}) - \tau_{\text{null}} \;=\; -\,s\,|s| \tag{42}$$

identically, for timelike, null and spacelike pairs alike. Requiring the tensor projection to reproduce the scalar one *on its trace*, pair by pair, is then one equation in one unknown,

$$W(s)\,s \;=\; T(\{a,b\}) - \tau_{\text{null}} \;=\; -\,s\,|s| \qquad\Longrightarrow\qquad W(s) \;=\; -\,|s|, \tag{43}$$

with no free function and no free scalar beyond the $8\pi G$ unit conversion (39) already carries.

One thing must be added, and nothing is read off it. With (43) and no range restriction the long pairs dominate and the answer drifts with the size of the region sampled, which a density must not do. PST supplies exactly one candidate scale for the restriction, the coherence length d_0, imposed on the invariant at $|s| \lesssim d_0^2$; once the window exceeds it the answer converges, to the same value for four different cutoff *shapes* and for cutoff scales differing by a factor of four (Computation 158). The cutoff is needed for convergence; no number comes out of it.

It is not a range restriction, and calling it one would mislead. The set $|s| \lesssim d_0^2$ is a neighbourhood of the *light cone*, which is not compact: a pair separated by a great distance still satisfies it, provided the separation is nearly null. Measured on a scattered cloud in a box of half-width 40, 98% of the surviving pairs are separated spatially by more than seven correlation lengths, with a median of 31 (Computation 159). That is inert for the vacuum and for radiation, and decisive for everything else, which is why the next subsection exists.

What that forces is radiation. By construction the trace of $\Pi_{\mu\nu}$ is the sum of $T(\{a,b\}) - \tau_{\text{null}}$ over pairs, and that vanishes in the large-window limit: the null threshold is the balance point of the pair-tension distribution. So the two

coefficients of (41) satisfy $\rho = 3p$, and

$$w \;=\; \tfrac{1}{3} \tag{44}$$

exactly, with $(\rho - 3p)/\rho$ shrinking from -1.9×10^{-1} to -5×10^{-6} as the window grows. It is also independent of the configuration: uniform, radially graded, Gaussian in space, Gaussian in space and time, and clustered into ten clumps all return w between 0.3333 and 0.3338. The pair-tension projection is *conformal matter*, and that is forced rather than chosen, which is what a pair structure carrying no scale of its own ought to give.

7.4.5. What is not derived

Two items were carried here. Computation 162 and Computation 163 have moved both, and what is left over is a third thing that was already named elsewhere.

First, conservation, and it now has two answers rather than none. $\nabla^\mu \Pi_{\mu\nu} = 0$ is what the field equations require, and the direct route to it does not close. Read covariantly a pair has no separation *vector*; the stand-in is Synge's world function derivative σ_μ, with $s = \sigma^\mu \sigma_\mu = 2\sigma$. Two identities are exact, that one and its gradient $\sigma^\mu \sigma_{\mu\nu} = \sigma_\nu$, and together they collapse the per-pair divergence to

$$\nabla^\mu \big[f(\sigma)\, \sigma_\mu \sigma_\nu \big] \;=\; \big[2\sigma f'(\sigma) \;+\; f(\sigma)\, \Box\sigma \;+\; f(\sigma) \big]\, \sigma_\nu, \tag{45}$$

so *all* the curvature dependence sits in one scalar, $\Box\sigma$, whose departure from its flat value 4 is $-\frac{1}{3} R_{\alpha\beta} \sigma^\alpha \sigma^\beta$ to leading order. On constant curvature that is exact in closed form, $\Box\sigma = 1 + 3(d/L)\cot(d/L)$, and with the derived weight (43) the relative shift is $-\frac{1}{7}(d/L)^2$. Since the pair sum is cut at $|s| \lesssim d_0^2$, the separation obeys $d \lesssim d_0 = 2.2$ am and the violation is minute everywhere the theory is used: 6×10^{-43} at a solar-mass horizon, 6×10^{-45} at a neutron star's surface, 8×10^{-59} at the Earth's, 5×10^{-88} at the cosmological curvature scale (on the data's bound $d_0 \leq 2.3$ nm each is eighteen orders larger and

still minute; the four values are Computation 162's table re-evaluated at the derived d_0). *That is the same parameter as* (66), where $(kd_0)^2 = 1.9 \times 10^{-42}$ decides whether the projected equations reach a horizon at all, so conservation can never be the dominant error. That much is unconditional.

And the exact answer, which costs an assumption rather than an approximation. §7.5 identifies $\Pi_{\mu\nu}$ as the conformal part of $T_{\mu\nu}^{(\text{fields})}$, and diffeomorphism invariance of S_{matter} makes that tensor conserved on shell. Splitting it as $T_{\mu\nu} = \Pi_{\mu\nu} + \frac{1}{4} g_{\mu\nu} T$ with Π traceless gives one line,

$$\nabla^{\mu} \Pi_{\mu\nu} \;=\; -\tfrac{1}{4}\, \nabla_{\nu} T, \tag{46}$$

exactly, on any background: the traceless part of a conserved tensor is itself conserved precisely when the trace is covariantly constant. Computation 158 showed the trace of $\Pi_{\mu\nu}$ *vanishes*, for every configuration tried, because the null threshold is the balance point of the pair-tension distribution. So the right-hand side is zero identically and $\Pi_{\mu\nu}$ is conserved *exactly*, with no $(d_0/L)^2$ residue at all, in the sector it actually describes. Computation 163 verifies (46) on a Friedmann background with a massive scalar, where the trace is not constant and both sides are nonzero, and confirms that the electromagnetic stress tensor is traceless for an arbitrary field strength, as the worked instance of a source in the class where the right-hand side vanishes.

So the construction does not need an action of its own, which is worth saying because it was carried here as though it did. What it needs is the identification, which Computation 199 leaves without obstruction (below), not the conservation.

Second, the divergence-free ambiguity, which is smaller than its name. A stress tensor assembled from pair contributions is fixed only up to the double divergence of a tensor with the algebraic symmetries of Riemann, the standard Irving–Kirkwood and Hardy freedom. For the one-scalar family $H_{\mu\alpha\nu\beta} = f\,(g_{\mu\nu} g_{\alpha\beta} - g_{\mu\beta} g_{\alpha\nu})$ the shift is $\Delta_{\mu\nu} = g_{\mu\nu} \Box f - \nabla_{\mu} \nabla_{\nu} f$, and Computation 163

computes three things about it exactly: it is symmetric, so it is an admissible regauging; its trace is $3\Box f$; and its divergence is $-R_{\nu\sigma}\nabla^{\sigma} f$, which is identically zero in flat space and curvature-suppressed otherwise, the same shape as the first item.

And tracelessness closes most of it. Keeping $\Pi_{\mu\nu}$ traceless is not a convention here: Computation 158 forces it on every balanced cloud, and out of balance the identification concerns the conformal part of the sum, which is traceless by definition (Computation 199). By the trace above that restricts the gauge to $\Box f = 0$, and a harmonic gauge function that is localised is identically zero, by uniqueness for the Dirichlet problem on the Riemannian section and by the domain of dependence in Lorentzian signature, both checked numerically. So no regauging confined to a region exists at all, and what survives is a global harmonic function, which changes no integrated quantity. The tensor is determined up to that, which is a long way from being a representative of an arbitrary class.

What is left, then, is the conformal identification, and it is no longer open. Both results above lean on $\Pi_{\mu\nu}$ being the conformal part of $T_{\mu\nu}^{(\mathrm{fields})}$, and §7.5 shows that nothing obstructs it: the conformal part of the pair sum reaches every metric-traceless tensor that obeys the null energy condition (Computation 199), and its magnitude is the state's (Computation 196). The item was carried in three places, then in one, and closed there.

One correction comes with all this, because the earlier wording invited the wrong reading. With the derived weight the per-pair divergence is *not* zero even in flat space; it is $-14\sigma\,\sigma_{\nu}$. "Checked in flat space" was therefore never a statement about W. It is a statement about the sum over pairs and the bond kernel, and Computation 162 measures it as such: the residual falls with the pair count, at a rate consistent with a finite-sample effect rather than a systematic.

The cosmological constant was never among these, and its absence from this

sector is by design rather than by default. Computation 155 showed that $w = -1$ is exactly the scalar form $\Pi\, g_{\mu\nu}$, which §16.5 derives separately as the vacuum term. The two sectors do different jobs and add: $T_{\mu\nu} = \Lambda\, g_{\mu\nu} + \Pi_{\mu\nu}$, the first with $w = -1$ and the second traceless.

7.4.6. What the matter sector can and cannot be

Real matter is not a scattered cloud of realisation points; it is made of *worldlines*. Separating the pairs that lie along one particle's own worldline from the pairs between different particles turns the equation of state into something recognisable, and then rules something out.

The worldline sector is exactly kinetic. Taking only the intra-worldline pairs, the projection gives

$$w \;=\; \frac{\beta^2}{3}, \tag{47}$$

exactly, verified to 5×10^{-15} over nine speeds (Computation 159). At $\beta = 0$ the pressure is zero to every printed digit, which is dust; at $\beta \to 1$ it is $1/3$. Equation (47) is the textbook kinetic equation of state of a gas, and nothing was put in to obtain it: it follows from the weight (43), which was itself forced.

A photon contributes nothing from its own worldline, and that is why radiation is $1/3$. A massless particle's worldline is null, so the invariant along it is $s = 0$ exactly, and the derived weight $W(s) = -|s|$ vanishes *identically* on it. Measured, the summed weight over the intra-worldline pairs of null worldlines sits fourteen orders below the same configuration at $\beta = 0$, consistent with round-off in $s = \Delta \cdot \Delta$. So the radiation value is not a statistical limit or a high-temperature approximation. It is forced by the weight vanishing on null separations, and the separation that is null is the one τ_{null} defines: the same cut that fixes the signature and gives the light cone. The photon enters the stress tensor exactly where it enters the geometry, through the same object.

But the worldlines are a vanishing share, and that closes the matter question. Writing f for the worldline share of the energy density, the two sectors mix as $w = \big[1 - f(1 - \beta^2)\big]/3$, so dust needs $f \to 1$ as well as $\beta \to 0$. It does not

get there. Because $|s| \lesssim d_0^2$ selects a neighbourhood of the light cone rather than a compact region, the near-null background never dilutes: sweeping the density over fifty times the correlation length leaves f between 0.01 and 0.05 throughout. The bond kernel of (40) is the obvious remaining candidate and it cannot help either, because it is *normalised*: $\int B \, \mathrm{d}^4x = 1$, so the volume-averaged density is identically the integrated sum, whatever B is. It decides where the stress sits, not what the equation of state is.

So the projection is **strictly conformal for homogeneous matter**, and dust is excluded rather than merely unfound. The rest mass of ordinary matter does not enter the stress tensor through pair separations at all. That is a real limit, stated as one, and what it does is move the question rather than leave it vague: PST has mass elsewhere, in the elementary mass $m_0 = 4\pi^3 r_0^2 d_0/(3c^2)$ of §9.5 and in the Higgs sector, so what survives is how the matter action rather than the pair geometry carries it into $T_{\mu\nu}$.

7.5. The field equations as a conservation statement

Einstein's field equations then read:

$$G_{\mu\nu} \;=\; 8\pi G \, T_{\mu\nu}, \qquad T_{\mu\nu} \;=\; \Lambda \, g_{\mu\nu} \;+\; T_{\mu\nu}^{(\mathrm{fields})} \tag{48}$$

Here Λ denotes the vacuum energy density; the constant of the Friedmann equations in curvature units is $8\pi G\Lambda$, which is how §16.5 and Appendix A quote the observed value. The first term is the vacuum: conservation forces the scalar projection $\Pi(T(C))$ of (39) to be constant, so it is exactly $\Lambda g_{\mu\nu}$ and can be nothing else (Computation 155). The second is the stress-energy of the projected matter fields, $T_{\mu\nu}^{(\mathrm{fields})} = -2(-g)^{-1/2}\,\delta S_{\mathrm{matter}}/\delta g^{\mu\nu}$, the same object (51) obtains by varying (50), whose content is the Standard Model that Chapters 12 and 13 derive.

Where the pair-tension projection sits in that, and where it does not. $\Pi_{\mu\nu}$ of (40) is not a third term and is not the whole of $T_{\mu\nu}^{(\mathrm{fields})}$ either. It is that tensor's *conformal part*: Computation 158 and Computation 159 show it is

traceless for homogeneous matter and reproduces the radiation sector exactly, $w = 1/3$, while carrying no rest mass at all. Rest mass reaches $T_{\mu\nu}$ through the matter action, by way of the Higgs and Yukawa structure of Chapter 14, and not through pair separations. Writing $T_{\mu\nu} = \Pi(T(C))\, g_{\mu\nu}$, or $\Lambda g_{\mu\nu} + \Pi_{\mu\nu}$, would say that the universe contains a cosmological constant and radiation and nothing besides: no matter era, no structure formation. It contains those because $T^{(\text{fields})}_{\mu\nu}$ does, for the ordinary reason that a massive field in its non-relativistic regime is pressureless.

The conformal identification, and it now has a criterion rather than a caveat. PST's ambition here is that geometry and matter are coprojections of one tension, and the operator Π is what carries that claim. Computation 159 showed Π delivers the conformal part and cannot deliver the rest. Computation 164 goes further and says exactly which conformal tensors it delivers, in both directions, which is a better position than the caveat it replaces.

The criterion. Every term of (40) is a rank-one dyad $\Delta_\mu \Delta_\nu$ carrying a weight $W(s) = -|s|$ that never changes sign, times a bond kernel that is a density. A conic sum of semidefinite matrices with same-signed coefficients is semidefinite, so $\Pi_{\mu\nu}$ is *always* semidefinite, and conversely every semidefinite target is a sum of dyads. **The reachable set is exactly the semidefinite metric-traceless cone**, sharp in both directions, and confirmed over six ensembles, an isotropic Gaussian cloud and five deliberately awkward ones: flattened to a plane, stretched along time, split into two clumps, laid on a light cone, collapsed onto a line. Not one negative eigenvalue. That is the pair sum itself; what the identification concerns is its conformal part, which reaches further, and the rest of this section says how far.

What that class contains is everything this book rests on here. Isotropic radiation, $\mathrm{diag}(\rho, \rho/3, \rho/3, \rho/3)$, is positive definite. A null plane wave, $\Phi\, k_\mu k_\nu$, is rank one and semidefinite. So Computation 158's $w = 1/3$, Computation 159's kinetic law, Computation 160's polarisation count and Stefan–Boltzmann constant and Computation 161's Planck spectrum all live inside the reachable

class, and none of them is touched by what follows.

And one exclusion stops being a design choice. $\Lambda g_{\mu\nu}$ is indefinite, so the pair sum could not have produced it whatever one wanted. The separation of the vacuum term into its own sector, which §16.5 argues on physical grounds, is *forced* by the structure as well.

What lies outside, stated plainly. A static magnetic field has $T_{\mu\nu} = \mathrm{diag}(u, u, u, -u)$ with $u = B^2/2$: metric-traceless, and so conformal by this book's own classification, but *indefinite*, the negative eigenvalue being the tension along the field lines. An electric field is the same. No pair sum equals it, and that was for a time the first structural item. The reason is physical rather than technical: a bond under a single-signed weight either pushes or pulls throughout, and a magnetic field does both at once, tension along the lines and pressure across them.

The one freedom left after Computation 163, a constant metric-traceless shift, does restore semidefiniteness, so what fails is the identification *as stated* rather than a gauged version of it. The price is that the constant is universal, must dominate the largest tension anywhere, and is not fixed by anything in the construction, since ρ is a quotient map onto M and leaves no empty region to pin it against.

And the obstruction does not reach the identification (Computation 199). The identification is a statement about conformal parts, and this book defines the conformal part of a tensor as the tensor minus a quarter of its metric trace times the metric. Applied to the pair sum term by term, each dyad becomes $\Delta_\mu \Delta_\nu - \frac{1}{4}\, s\, g_{\mu\nu}$, traceless identically and *indefinite*: a spacelike bond carries pressure along itself and tension across itself. Two families of bonds in the plane transverse to a direction sum to exactly $\mathrm{diag}(u, u, u, -u)$ with the field along that direction, and the electric tensor is the same construction in the other transverse plane; both are reached to machine precision from explicit clouds. The reachable set is sharp in both directions again, and larger: the

closed cone the traceless dyads span is *exactly the metric-traceless tensors obeying the null energy condition*, $T_{\mu\nu}k^\mu k^\nu \geq 0$ for every null k. Inclusion is one line, $(\Delta \cdot k)^2 \geq 0$ with the metric term killed by k's nullity; equality is a duality argument, the dual cone being the positive semidefinite metric-traceless matrices, whose extreme rays are rank one by Pataki's bound on spectrahedra and therefore null dyads, the rank-two case checked by hand on the two null directions of its plane. Every classical field obeys the condition, so every conformal part a field of this book produces is reachable, and what is excluded is exactly a null-energy violation: a traceless tensor with negative energy density along some null direction, which nothing here contains. On every balanced cloud the conformal part is the pair sum itself, the trace being Computation 158's balance scalar and vanishing there, so Computation 158 to Computation 161 are untouched; out of balance the sum carries a trace the identification does not use, and Computation 159 had already met it: the worldline sector of the pair sum is exactly kinetic, $w = \beta^2/3$, dust at rest, and it carries a few per cent of the density for homogeneous matter, which is the share by which the projection there is conformal rather than exactly traceless. The forced weight is load-bearing in a new way: a weight changing sign across the light cone would flip the spacelike dyads and the cone would reach null-energy violations, so $W(s) = -|s|$ is what confines the pair sum to physical matter. The universal constant shift above is no longer needed, and the magnetic tensor sits on the boundary of the cone, the condition saturating along the field lines, as a static field should.

And a second thing is free, which belongs beside the first. The weight is forced and the bond kernel normalised, so the only remaining freedom in $\Pi_{\mu\nu}$'s *magnitude* is how many pairs there are, and the theory does not supply that: d_0 is a coherence length, which §15.3.1 says is not the microscopic grain, and ℓ_* enters as a momentum cutoff on field modes. Reading the spacing off one rather than the other moves the implied density by $(d_0/\ell_*)^6 \approx 10^{61}$ at the Casimir bound (Computation 167; 9×10^6 at the derived d_0, Computation 201). Nothing claimed here reads that scale, which is why nothing claimed here

is affected: every result the pair sum has produced is a ratio, a shape or a count, and Computation 167 verifies the invariance rather than assuming it. One consequence deserves stating plainly, because a reader will reach for it: there is *no* vacuum-energy catastrophe in this sector. That problem needs a computed scale to disagree with measurement, and this construction does not produce one.

The density is not missing, though, and saying so changes the shape of the item (Computation 179). Computation 167 searched among the theory's *lengths* and found neither to be a spacing, which is right as far as it goes; A1 names one outright. §18.1 introduces a_n as "the substrate site spacing, distinct from the fixed coherence length d_0", with $|D_n|a_n^4 \to V_M$, $a_n \to 0$ and local node density $\lambda_n = 1/a_n^4 \to \infty$, and it scales the *same* substrate: (40) sums over pairs of D through ρ. So the density is supplied and it is infinite. With the pair cutoff the book itself uses, $|s| \lesssim d_0^2$ at the *fixed* coherence length (Computation 159), a cutoff that does not scale against a density that does gives $\Pi_{\mu\nu} \sim |D|^2 \sim a_n^{-8}$, measured, and Computation 182 confirms that the exponent survives the bond kernel's localisation rather than being an artefact of summing globally.

Which D that is, is not a choice (Computation 182). It reads as one, between A1's sequence and one physical substrate of finite spacing, and the second is not available. D is finite at every stage, by four independent commitments of this book: A1's own $|D_n| \to \infty$, the Bernoulli product $\bigotimes_{a\in D}$, the exact enumeration over $2^{|D|}$ configurations of §10.3.2.5, and the balance probability $\binom{|D|}{|D|/2}2^{-|D|}$ of §3.5. A finite set has no uncountable quotient, so the quotient $M = D/\!\sim_\rho$ of §4.10 is finite at every stage and is a smooth four-manifold at none of them, while (46) takes a covariant derivative of $\Pi_{\mu\nu}$ on a smooth M and (41) decomposes it against a metric. So D in (40) is A1's sequence, and $M = \mathbb{R} \times S^3$ is the limit of the finite quotients rather than the quotient of any one substrate. Computation 167's arithmetic is reproduced exactly and is not in question; what falls is its reading of the symbol, and Computation 179 recorded that as an unstated premise where the sharper statement is that it

was an unavailable one.

And the forced normalisation does not buy a finite density, which is the correction this makes to Computation 179. Dividing by the pair count is forced and it converges. But (40) is *dimensionless*: $W(s)$ carries L^2, the unnormalised dyad another L^2, and a bond kernel normalised to unit integral carries L^{-4}, so the bare sum is a pure count, which is why it goes as $|D|^2$. A stress tensor is an energy density, and the only scale the discrete data carries is a_n, so restoring the dimensions costs exactly one factor a_n^{-4}, which is λ_n itself. Measured, the restored object diverges as *one* power of the node density, fitted at 1.011 against controls at 0.049 and 2.049 (Computation 182).

So M7's missing quantity is named, and there is exactly one of it: the physical spacing a_* at which A1's idealised limit is truncated, with $\Pi_{\mu\nu} \sim 1/a_*^4$. That is the same object a lattice spacing is, and it puts this construction in the ordinary position of a lattice one: the continuum limit does the geometry, the finite spacing sets the magnitude, and no result that is a ratio, a shape or a count feels it. Computation 167's finding that neither d_0 nor ℓ_* is such a spacing stands, which is why the item does *not* close; what changes is that the gap is one truncation scale of a familiar kind rather than a density the theory was said not to supply. The vacuum-energy conclusion above is unaffected, since it needs only that no absolute scale is *computed*, and a_* is named rather than computed.

And a_* is not a constant of the theory (Computation 196). On the Friedmann geometry of Chapter 16, (46) with the trace zero makes a homogeneous traceless $\Pi_{\mu\nu}$ fall as $a(t)^{-4}$, for every spatial curvature; a fixed physical a_* is the case $\rho = \text{const}$, whose divergence is $4H\rho$, so the truncation spacing is comoving. Under the identification this section makes, the traceless homogeneous part of $T_{\mu\nu}^{(\text{fields})}$ is the radiation and nothing else ($\rho - 3p = 0$ iff $w = 1/3$), so ${\Pi^0}_0 = \rho_r$ and $a_* = (\hbar c/\rho_r)^{1/4} = 0.82$ mm today, which is 0.97 times $\hbar c/kT_{\text{CMB}}$: the thermal wavelength of the radiation bath, scaling as $a(t)$. It passes through the Casimir bound on d_0 at $kT = 83$ eV, through the derived d_0 at

$kT \approx 40$ GeV (Computation 201) and through $0.4\,\ell_*$ at $kT = M_*$, which is Computation 167's negative result as a function of time. So the magnitude of $\Pi_{\mu\nu}$ is supplied by the state, as the magnitude of $T_{\mu\nu}$ is in general relativity; with Computation 199 removing the cone restriction the identification has no gap left and is carried in Appendix C.6; the vacuum-energy conclusion is untouched because ρ_r is read, not computed.

This is the same unstated choice Computation 177 found in P3, in a second construction. There (55) was a bare sum over D while (57) divides by $|D|$, so a constant coefficient diverged under A1 until Computation 177 showed c must carry a $1/|D|$; here (40) is a bare sum over *pairs* of D carrying no normalisation at all. *Every bare sum over D is sensitive to which reading is meant*, and that is worth checking in any further construction of this kind.

A second doubling sits underneath it and is worth stating separately, because it is about a *definition* rather than a sum. This book gives the realisation map twice and the two are not of the same type. §4.10 makes ρ a *surjective quotient* whose target is derived, "M is the quotient set $D/\!\sim_\rho$", with a spacetime point defined to *be* a fibre, and it spends the surrounding paragraphs establishing that M is not a container given in advance. §18.1 then writes $\rho_n : D_n \to M = \mathbb{R} \times S^3$ with image "asymptotically equidistributed", which for finite D_n cannot be surjective and takes M as given; and the assumption list of §7.8 says as much outright, that ρ "presumes the target manifold M is smooth and finite-dimensional". The reconciliation is the one above: the quotient is the stage-wise statement and the smooth M is the limit. The fibre structure, which an earlier printing named here as unsaid, is settled on the theory's own measure (Computation 212): a fibre is a clique of pairwise-null properties, a set of directions in V_7 with one pairwise inner product, so it holds at most seven properties (the Gram rank), and under the G_2-invariant law on directions an exact null coincidence has probability zero, so fibres are singletons almost surely. $|D_n|a_n^4 \to V_M$ counts *properties* against *volume*, and A1 reads $|D_n|$ as the point count; the exponent above is insensitive to the fibre size in any case, and §4.10's identification of a point *with* a bundle of

properties holds with a bundle of size one, at most seven (Computation 182).

So the honest statement is this, and it is a stronger one than before. $\Pi_{\mu\nu}$ *is* the conformal part of $T_{\mu\nu}^{(\text{fields})}$: on the semidefinite sector by the characterisation above, which contains every result claimed here, and beyond it by Computation 199, which removed the cone restriction, so the identification has no gap left and no universal constant to supply. A framework that can say exactly which part of $T_{\mu\nu}$ it derives, and from what, is in a better position than one that takes the whole of $T_{\mu\nu}$ as input and never reaches the question. General relativity is the second kind.

7.5.1. Why both routes are kept

The question this invites is whether the projection is then redundant. If the matter action supplies $T_{\mu\nu}^{(\text{fields})}$, and supplies exactly the part Π cannot, what work is left for Π to do? Three things, and the matter action does none of them.

It supplies Λ, and the matter action cannot. The vacuum energy of S_{matter} is the Standard Model's, with the hierarchy problem §16.5 says plainly PST inherits. What the projection adds is a term of a kind no field on the manifold can produce: $\rho_{\text{inst}} = j_\tau \, \Delta\mathcal{F}/c^2$ of (221) does not dilute, because the configuration space it draws on is precausal and has no spatial extent to stretch. That fixes the *sign* of Λ, positive and de Sitter, from the direction of every threshold crossing, and fixes $w = -1$ as a theorem conditional on the pre-spatial framing (§9.1) rather than a fit (Computation 137, Computation 155). The magnitude stays $T(C)$-contingent and is not claimed. Without the projection, Λ is once again a number to be fitted, which is the position PST set out to improve on.

It makes the radiation sector a derivation rather than an input. In the matter action, $w = 1/3$ holds because a massless field with a conformal action was written down; the answer was already in the premise. In the projection it holds because the derived weight $W(s) = -|s|$ of (43) vanishes identically on null separations, and the separation that is null is the one τ_{null} defines in (30).

§8.3.4 carries that further: the same threshold fixes the two polarisations, the Stefan–Boltzmann law and its coefficient. Those are facts about light obtained from the object that fixes the signature, and S_{matter} on its own reaches none of them.

It is where the unification claim has content, and it is checkable. A formulation with the matter action alone is General Relativity with a prior account of the manifold: geometry from the substrate, matter handed in from outside. That is less than PST claims. And because $\Pi_{\mu\nu}$ is identified with the conformal part of $T^{(\text{fields})}_{\mu\nu}$ rather than with all of it, the identification is a statement about a sector that is independently known, so it can fail: the traceless part of the matter stress tensor has to equal the pair-tension sum (40). A formulation without the projection has nothing of that kind to fail.

So the two are not rival suppliers of the same tensor, one of which has been made redundant. The projection's job has changed: from carrying the whole right-hand side, which Computation 159 rules out, to carrying the vacuum term, carrying the conformal sector, and constraining the rest.

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.

7.6. The F–S dichotomy: free energy versus action

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

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

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

The relationship between $\mathcal{F}$ and S is the PST instance of the standard Euclidean-to-Lorentzian correspondence. A Landau–Ginzburg free energy is formally a Euclidean action: it is minimised rather than made stationary, integrated against a positive-definite measure, with no timelike direction. The

standard route from Euclidean to Lorentzian is Wick rotation: $t \to -i\tau$ reverses the sign of the time-kinetic term and changes the signature from $(+,+,+,+)$ to $(+,-,-,-)$. In PST, sublimation *is* the Wick rotation. The moment the $(+,-,-,-)$ Lorentzian signature emerges from the functor Φ is precisely the moment $\mathcal{F}$ ceases to be the relevant functional and S takes over. This is not an embarrassing gap in the formalism. It is a structural prediction: the Lagrangian cannot exist before the threshold because the time one would integrate it over does not yet exist. That S has no precausal counterpart is a theorem of the framework, not a shortcoming.

F–S Dichotomy. $\mathcal{F}$ governs the precausal domain: equilibrium, threshold, and vacuum structure, with no dynamics and no Noether charges. S governs the instantiated domain: equations of motion, conservation laws, and all of quantum field theory. The passage from $\mathcal{F}$ to S is sublimation; it is not reversible; and nothing on the S-side can be traced back to the precausal level without crossing the threshold. Every subsequent use of Noether's theorem, every conserved charge, and every field-theoretic result in this book belongs exclusively to S.

7.7. The projected action and the Einstein–Hilbert term

Newton's constant from the spectral action. Newton's constant is the gravitational coefficient of the spectral action $\mathrm{Tr}\, f(D/\Lambda)$ [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 [155] and Higgs terms at order Λ^0. The Seeley–DeWitt heat-kernel computation [57] carried out in this section 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 Λ the spectral-action cutoff, which G itself fixes (§8.5) and which is neither the coherence length d_0 nor the grain ℓ_* (§15.3.1). 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 (48) were presented above as a conservation statement,

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

The potential functional $\mathcal{F}(\psi, C)$ of equation (77) (stated in full later) contains three terms. The first two, $-\varepsilon\,|\psi|^2$ and $b|\psi|^4$, determine the equilibrium structure of the order parameter. The third, $c\,|\nabla_{\mathcal{C}}\psi|^2$, is the gradient term: it penalises variation of the order parameter across configuration space and controls the energetic cost of excitations above the vacuum. Under the projection Φ, this term acquires a precise geometric interpretation.

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

$$\Pi\big(c\,|\nabla_{\mathcal{C}}\psi|^2\big)\,\big|_{\psi=r_0} \;\longrightarrow\; \frac{r_0^2\,c}{16\pi G}\int_M R\,\sqrt{-g}\;\mathrm{d}^4x \;+\; \mathcal{O}\big(R_{\mu\nu}R^{\mu\nu},\,R^2\big) \tag{49}$$

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 (72) (derived later, §8.9). 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:

$$S[g] \;=\; \frac{1}{16\pi G}\int_M R\,\sqrt{-g}\;\mathrm{d}^4x \;+\; S_{\mathrm{matter}}[g,\,\mathrm{fields}] \tag{50}$$

This is the **Einstein–Hilbert action**. The step from the gradient term in $\mathcal{F}$ to equation (50) is structural. The functional form of the projection

(that the leading curvature term is $\int R\sqrt{-g}\,\mathrm{d}^4x$ and not a higher-curvature combination) is derived in §7.8 via the Seeley–DeWitt heat-kernel expansion. That resolves the projection-coefficient question: the structure of the R term is established here, and its numerical coefficient $\mathcal{N}$ is identified as the finite-internal-factor trace in §7.1. The downstream Noether–Lovelock argument [22] is standard and unambiguous given equation (50). Varying equation (50) with respect to $g^{\mu\nu}$ yields, by a standard calculation:

$$\frac{\delta S}{\delta g^{\mu\nu}} \;=\; 0 \qquad \Longrightarrow \qquad G_{\mu\nu} \;=\; 8\pi G\, T_{\mu\nu} \tag{51}$$

where $T_{\mu\nu} = -2/\sqrt{-g}\cdot\delta S_{\mathrm{matter}}/\delta g^{\mu\nu}$ 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 (48) is thereby derived, not asserted.

7.8. Seeley–DeWitt argument for the projected action

The previous section identified the projection coefficient $\mathcal{N}$ in equation (49) as the trace over the finite internal factor. This section shows that the *functional form* of the projection (the leading curvature term $\int R\sqrt{-g}\,\mathrm{d}^4x$ announced at equation (50)) is forced by a universal result from differential geometry: the Seeley–DeWitt heat-kernel expansion [57, 59]. The numerical coefficient (which fixes G in terms of r_0, c, and d_0) is delivered by the spectral-action moments of Computation 5, $1/(16\pi G) = (f_2\Lambda^2/\pi^2)\cdot\kappa_{\mathrm{grav}}\cdot N_{\mathrm{dof}}$, with the matched cutoff Λ and the fermionic multiplicity N_{dof} as the calibration inputs. Pinning G to its observed value via this relation fixes one cutoff combination at the Planck scale (consistent with the $M_* = 4\pi m_h\sqrt{2/3}$ matched scaling at the electroweak sector); the numerical coefficient is therefore a *calibration* of the cutoff identification rather than an open derivation, on the same footing as the coherence length d_0 (structurally a condensate quantity, 2.2 am under the canonical normalisation, Computation 201, and bounded by 2.3 nm from the data via (217)).

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

$$\Pi(A)[x] \;=\; \int_{\mathcal{C}} A[\psi]\, K(\mathcal{C}, x)\, \mathrm{d}\mu(\mathcal{C}) \tag{52}$$

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

$$|\nabla_{\mathcal{C}}\psi|^2 \;=\; r_0^2\, |\nabla_{\mathcal{C}}\theta|^2 \;+\; |\nabla_{\mathcal{C}}\sigma|^2 \;+\; \mathcal{O}(\sigma^2) \tag{53}$$

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 (Chapters 13 and 14). 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:

$$\Pi\big(c\, |\nabla_{\mathcal{C}}\sigma|^2\big)\,\big|_{\sigma=0} = \; c\, r_0^2 \int_M \widetilde{K}(x, x;\, d_0^2)\, \sqrt{-g}\; \mathrm{d}^4 x \tag{54}$$

where $\widetilde{K}(x, y;\, t)$ 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 $\widetilde{K}$ to be, to leading order, the heat kernel of a *Laplace-type* operator $P = -(g^{\mu\nu}\partial_\mu\partial_\nu + \cdots)$ whose principal symbol is the spacetime metric g. This is not a consequence of smoothness, positivity, symmetry, and localisation: a normalised 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 localised kernels.

The Laplace-type condition is derived from P1–P3 under A1–A7 via Mosco convergence (§7.8), and stated separately in §1.7 for reference. It runs as follows. $\widetilde{K}$ 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 $\widetilde{K}$; 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 $\int R\sqrt{-g}\,\mathrm{d}^4x$.

Taken as an independent postulate, the Laplace-type projection condition 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 as 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 Laplace-type projection condition is derived by an argument using Mosco convergence of Dirichlet forms [98]. Define the *Boolean Dirichlet form* on $L^2(\mathcal{P}(D), \mu)$ by

$$E^B(u, v) \;=\; \sum_{a \in D} \int_{\mathcal{P}(D)} (\partial_a^B u)(C)\,(\partial_a^B v)(C)\,\mathrm{d}\mu(C), \tag{55}$$

where ∂_a^B is the Boolean difference operator of equation (20) and the generator of E^B is the Boolean Laplacian L^B. For smooth functions $f, h \in C^\infty(M)$ pulled back via $\Pi^\dagger$, the isotropy of the Bernoulli measure μ forces the covariance

matrix of the displacement $\rho(a)^\mu - x^\mu$ to be proportional to $g^{\mu\nu}(x)$:

$$\frac{1}{|D|}\sum_{a\in D}\big(\rho(a)^\mu - x^\mu\big)\big(\rho(a)^\nu - x^\nu\big) \;\longrightarrow\; \tfrac{d_0^2}{4}\, g^{\mu\nu}(x) \tag{56}$$

(no preferred direction exists in the substrate, so the only $O(4)$-covariant rank-2 tensor built from g is $g^{\mu\nu}$ itself). The isotropy here is Euclidean, and it fixes only the positive-definite content: the scale $\frac{d_0^2}{4}$ and the absence of a preferred spatial direction. A covariance is positive-semidefinite and cannot itself carry indefinite signature. The Lorentzian signature of g is therefore supplied separately by the τ_{null} sign cut ((30), Computations 2 and 124), and (56) is written in that post-signature metric. It follows that

$$\begin{aligned}\frac{1}{|D|}\,E^B\big(\Pi^\dagger(\Pi f),\,\Pi^\dagger(\Pi h)\big) \;&\xrightarrow[|D|\to\infty]{\text{Mosco}}\; \frac{d_0^2}{4\,V_M}\int_M g^{\mu\nu}\,\partial_\mu f\,\partial_\nu h\,\sqrt{-g}\,\mathrm{d}^4x \\ &=:\; \frac{d_0^2}{4\,V_M}\,E^{\text{Riem}}(f,h),\end{aligned} \tag{57}$$

where E^{Riem} is the Riemannian Dirichlet form whose generator is the Laplace–Beltrami operator Δ_g. *The $1/V_M$ is the same normalisation A2 carries, and it is not optional*: the left side is an expectation against the probability measure μ while E^{Riem} as written is an unnormalised integral, so the two cannot be equated without it, and Computation 177 measures the factor on an explicit lattice and recovers it to 0.3% across a factor 81 in volume. As with A1 and A2 the statement is read locally, V_M being infinite on $M = \mathbb{R}\times S^3$, so V_M is $\mathrm{vol}(K)$ on the compact K carrying the test function; equivalently, and more cleanly, the limit is $\frac{d_0^2}{4}\langle g^{\mu\nu}\partial_\mu f\,\partial_\nu h\rangle$ with $\langle\cdot\rangle$ the normalised average. By Mosco's theorem [98], this Γ-convergence of quadratic forms implies strong resolvent convergence of their generators:

$$\big(|D|^{-1}L^B + \lambda\big)^{-1}\Pi^\dagger \;\xrightarrow{|D|\to\infty}\; \Pi\big(\Delta_g + \lambda\big)^{-1} \qquad \forall\,\lambda > 0. \tag{58}$$

The generator of the limiting form is Δ_g, a Laplace-type operator with principal symbol $g^{\mu\nu}$, so $\Pi(L^B)$ is Laplace-type in the scaling limit: the

projection condition follows from P1 (the distinction primitive, with the Bernoulli measure derived in Chapter 2, §2.5), P2 (asymmetric tension) and P3 (the Landau–Ginzburg modal potential), 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 $\rho : D \to M$ 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 (30).

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 real K-theory (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 (§4.12); 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. A complementary physical mechanism supplies the same sign cut. The Mosco scaling limit by itself yields Euclidean signature in the $|D| \to \infty$ 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 [181] $O(1,3)$. 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 (§4.12) 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 $d_0 \to 0$ gives [57, 58, 59]:

$$\widetilde{K}(x,x;\,d_0^2) \;=\; \frac{\mathcal{N}}{(4\pi d_0^2)^2}\left[1 \;+\; \frac{d_0^2}{6}\,R(x) \;+\; d_0^4\,\mathcal{A}_2(x) \;+\; \mathcal{O}(d_0^6)\right] \qquad (59)$$

where $\mathcal{N}$ is the normalisation of the kernel (fixed by $\int_M \widetilde{K}\,\sqrt{-g}\,\mathrm{d}^4x = 1$ 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 $\mathcal{A}_2(x) = \frac{1}{180}(R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} - R_{\mu\nu}R^{\mu\nu} + 6\nabla^2 R) + \frac{1}{72}R^2$ 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 (59) into equation (54):

$$\begin{aligned}\Pi\big(c\,|\nabla_{\mathcal{C}}\sigma|^2\big)\,\big|_{\sigma=0} &= \frac{c\,r_0^2\,\mathcal{N}}{(4\pi d_0^2)^2}\int_M \sqrt{-g}\,\mathrm{d}^4x \\ &\quad + \frac{c\,r_0^2\,\mathcal{N}}{6\,(4\pi)^2\,d_0^2}\int_M R\,\sqrt{-g}\,\mathrm{d}^4x \;+\; \mathcal{O}(d_0^4\,R_{\mu\nu\rho\sigma}^2)\end{aligned} \tag{60}$$

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 (49) 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 (49).

The coefficient and Newton's constant. Comparing the coefficient of $\int R\sqrt{-g}\,\mathrm{d}^4x$ in equation (60) with $1/(16\pi G)$:

$$\frac{c\,r_0^2\,\mathcal{N}}{6\,(4\pi)^2\,d_0^2} \;=\; \frac{1}{16\pi G} \qquad\Longrightarrow\qquad G \;=\; \frac{6\,(4\pi)^2\,d_0^2}{16\pi\,c\,r_0^2\,\mathcal{N}} \;=\; \frac{6\pi\,d_0^2}{c\,r_0^2\,\mathcal{N}} \tag{61}$$

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

$$\mathcal{N} \;=\; \int_{\mathcal{C}} K(\mathcal{C}, x)\,\mathrm{d}\mu(\mathcal{C}) \tag{62}$$

evaluated at any point $x \in M$ (the result is independent of x by permutation invariance of μ). This defines the normalisation $\mathcal{N}$. In the spectral-triple reading, $\mathcal{N}$ is identified in §7.1 with the trace over the finite internal factor, and equation (61) then fixes G as a consistency relation among c, r_0, and d_0 (§8.5).

Summary. The Seeley–DeWitt argument establishes three things without additional assumption. First: the functional form of the gradient projection is $\int R\sqrt{-g}\,\mathrm{d}^4x$ 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 (61); $\mathcal{N}$ is identified in

§7.1 as the trace over the finite internal factor, and G is fixed as a consistency relation among (m_h, M_P, v).

7.9. Summary

This chapter derives energy and matter as projections of the same asymmetric tension that projects to geometry. It leans on four analytic results stated in §7.1 and proved later (three in Chapter 18, the fourth in §8.10). The stress-energy tensor is the tensor projection $\Pi_{\mu\nu}$ of the tension functional (§7.4): its scalar part is exactly the cosmological constant and its conformal part the projected matter fields, so Einstein's equations read as a conservation statement rather than a source equation (§7.5). The modal free energy $\mathcal{F}$ and the instantiated action S are distinct functionals with distinct roles; the gradient term of $\mathcal{F}$ projects to the Einstein–Hilbert term, whose functional form the Seeley–DeWitt expansion forces at leading curvature order with higher-curvature corrections suppressed by $(k\, d_0)^2$ (§7.8). Newton's constant enters as the trace over the finite internal factor and is fixed as a consistency relation among (m_h, M_P, v), which Chapter 8 takes up.

8. Gravity from the Projected Geometry

In this book gravity is not a force added on top of the world. It is the shape of the world. This chapter checks that the shape comes out right: space bends the way Einstein said, falling things behave the way Newton said, and the strength of gravity turns out to be a conversion between the substrate's units and laboratory units rather than a free dial somebody had to set.

8.1. 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 (30) and the definition of the Ricci scalar:

$$R\big(\Phi(C)\big) \;\propto\; \Pi\big(\nabla^2_{\mathcal{C}}\, T(C)\big) \tag{63}$$

Curvature is the projection of the configuration-space Laplacian of asymmetric tension onto the instantiated manifold. The Laplacian $\nabla^2_{\mathcal{C}}$ acts on the discrete power set $\mathcal{P}(D)$; its continuum realisation via Π is delivered by the Mosco-convergence theorem (C1 of §7.1) together with the binomial-to-Gaussian projection kernel of Computation 73 (developed in §15.6). 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 propa-

gating 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 the Laser Interferometer Gravitational-wave Observatory (LIGO) [45] is fully consistent with this picture; on PST's reading the waves *are* propagating tension gradients, though the observation establishes the waves, not the reading.

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

Equation (51) is the Einstein field equation. For the claim that PST derives general relativity (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 fields part of the projected stress-energy, $T_{\mu\nu}^{(\text{fields})}$ of equation (48), reduces to a non-relativistic matter density, $T_{00}^{(\text{fields})} \approx \rho(x)$ with all spatial components negligible; the scalar projection $\Pi(T(C))$ is the constant Λ term of that equation and drops out of the Newtonian limit. Working in harmonic gauge with $\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h$, the linearised field equations become:

$$\nabla^2 \bar{h}_{00} \;=\; 16\pi G\,\rho \tag{64}$$

Setting $\bar{h}_{00} = 4\phi$, where ϕ is the Newtonian gravitational potential ($g_{00} \approx 1{+}2\phi$ in the book's mostly-minus convention), equation (64) reduces to:

$$\nabla^2 \phi \;=\; 4\pi G\,\rho \tag{65}$$

Poisson's equation for Newtonian gravity is recovered exactly, with G playing the role of the unit-conversion factor between the projected matter density $\rho = T_{00}^{(\text{fields})}$ and the metric perturbation ϕ.

8.2.1. Post-Newtonian parameters and scalar-mode decoupling

At post-Newtonian order, the same field equations reproduce perihelion precession, gravitational light bending, and the Shapiro delay at the GR values $\gamma = \beta = 1$. The non-trivial step is to show that PST's scalar degrees of freedom from the modal-sublimation sector do not introduce a graviscalar admixture that would shift γ away from unity, the standard signature of scalar-tensor deviations from GR. Four real scalar degrees of freedom are present in principle:

1. *The physical Higgs h (the Landau–Ginzburg [LG] radial / amplitude mode).* The PST–Higgs identification of §15.2 fixes $\phi = \Pi(\psi)$ with the radial mode acquiring mass $m_h^2 = 4\varepsilon$, i.e. $m_h \simeq 125$ GeV. The corresponding Yukawa range $\lambda_h = \hbar/(m_h c) \simeq 1.6 \times 10^{-18}$ m is dozens of orders of magnitude below any parameterised post-Newtonian (PPN) distance scale (orbital, lunar-laser-ranging, light deflection at the solar limb). The Higgs is therefore effectively static at PPN scales and contributes no long-range scalar gravity.

2. *The three electroweak (EW) eaten Goldstones, θ among them.* Under the Standard Model (SM) symmetry breaking $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{em}}$ (§13.4), three angular dofs of the Higgs doublet are absorbed into the longitudinal polarisations of $W^{\pm}$ and Z, giving these gauge bosons masses $m_W \simeq 80.4$ GeV, $m_Z \simeq 91.2$ GeV. These dofs no longer exist as independent scalars; they are part of massive vector fields with Yukawa ranges again far below PPN scales.

   The substrate's phase is one of the three, and this list used to count it twice. With $\psi = H^0$ (§13.4, §13.3; the same field in three vocabularies), the LG phase θ *is* the phase of the neutral component, which the broken diagonal combination takes up: it is eaten by the Z (Computation 172, Computation 171). The first edition gave it a fourth item of its own and

called it Stueckelberg-eliminable under unbroken $\mathrm{U}(1)_{\mathrm{em}}$. That reason does not work, and the fact it appeals to is what breaks it: $Q = T_3 + Y/2$ annihilates the vacuum, which is exactly why $\mathrm{U}(1)_{\mathrm{em}}$ is unbroken, but Q annihilates H^0 as well, H^0 carrying charge zero, so a $\mathrm{U}(1)_{\mathrm{em}}$ transformation does not move θ at all and has nothing to gauge away. *What the item was used for is unaffected*, since eaten and gauge-fixed both remove θ from the infrared (IR); only the mechanism changes. The cosmological zero-mode constraint $J^0 = 0$ (§9.1) remains consistent, and is stated there with the charge it belongs to, the broken neutral direction's rather than the electromagnetic one.

Hence no propagating long-range scalar dof from the modal-sublimation sector survives in the IR. The counting is four real components of the complex doublet, three of them eaten and one left physical and massive: the Higgs at 125 GeV, the eaten Goldstones absorbed into $W^{\pm}$ and Z at 80–91 GeV, and *no massless scalar at all* (Computation 172). No scalar-tensor admixture to the Schwarzschild geometry arises, and $\gamma = \beta = 1$ follows from the projected equation (51) being the pure Einstein equation at PPN order.

The sole physical differences from standard GR therefore appear at scales near d_0 (the Casimir correction of Chapter 15) and in the ontological provenance of G and Λ, not in any of the classical PPN observables.

8.3. The strong-field limit: horizons, and why light cannot leave

The weak-field limit of the previous section is the easy direction: it has to reproduce Newton, and it does. The opposite limit is the one a reader arrives with, and it is where the tension-gradient picture is most exposed, because the two things it has to account for (a horizon, and light that cannot cross it) turn out to be one threshold seen twice.

8.3.1. A horizon is a threshold surface, not a trap

Black holes are the extreme limit of the gradient picture of §8.1: they correspond to configurations in which $\nabla^2_{\mathcal{C}} T(C) \to \infty$ at a point in configuration space. The event horizon is the projection of the set of configurations at which the tension gradient reaches the null-threshold condition, beyond which no precausal configuration can project outward.

Two things about that sentence need saying plainly, because the register the rest of this chapter uses is stronger than the one available here. It is an *identification*, not a theorem: nothing above derives that the null-threshold set is what a horizon *is*, and neither the horizon area nor the surface gravity nor any entropy is computed from the substrate side of it. And nothing in it is a force. The horizon is not a place where a pull becomes too strong to escape; it is the surface at which the condition that permits outward projection stops holding. Those are different kinds of statement, and only the second is PST's to make.

8.3.2. Why light cannot leave, when nothing is pulling on it

The expected answer is the Newtonian one: gravity pulls the light back. That answer is not available here, and it should not be, because it quietly gives light a mass for gravity to pull on. PST's photon takes no mass term, being the connection of the unbroken $Q = T_3 + Y/2$ (§13.4); that is the Standard Model's own unbroken-generator argument, which PST inherits rather than derives independently, and it is used here as corroboration rather than as a premise, since the argument below needs nothing from it.

PST's own answer needs one object, and it is an object the reader already has. τ_{null} is a condition on *pairwise* tension: a pair with $T(\{a, b\}) = \tau_{\text{null}}$ is null-separated, one above it spacelike, one below it timelike ((30), §4.10). That single cut has already done three jobs in this book. It supplies the Lorentzian signature, so it *is* the light cone rather than a structure laid on top of one. It sets the propagation speed of a disturbance in $T(C)$, which is why gravitational waves came out travelling at the speed of light in §8.1

without that speed being put in by hand. And, in the paragraph above, it defines the horizon.

So: *light is not held back; it runs exactly along* τ_{null}*, and a horizon is the surface at which that same* τ_{null} *closes off outward projection.* What fixes where light can go and what constitutes the horizon are the same object seen twice. No force acts on the photon anywhere in that account, and none is needed, which is the point, because a theory that had to supply one would owe the photon a mass it does not have.

8.3.3. Where the projected field equations stop

Two things were said above about the horizon identification, and their natural reading overshoots to a third. It is tempting to conclude that (51), derived in a low-energy expansion, simply does not reach the strong-field case. That is too broad, and the place to check it is the validity condition the expansion was given when it was stated: the higher-curvature terms at (49) are suppressed by $(k\,d_0)^2$ and negligible at all scales $\gg d_0$. That is a condition on *curvature*, and a horizon is a strong-*potential* region, not a strong-curvature one. The two come apart, and at a horizon they come apart by some forty-two orders of magnitude.

Take the curvature scale to be $k = K^{1/4}$, where $K = R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$; on the Schwarzschild geometry, whose Ricci tensor vanishes identically, $K = 12\,r_s^2/r^6$ is the only non-vanishing invariant, so at the horizon

$$(k\,d_0)^2\Big|_{r=r_s} \;=\; \sqrt{12}\left(\frac{d_0}{r_s}\right)^2 \;\approx\; 1.9\times10^{-42} \quad \text{for } M = M_\odot, \tag{66}$$

and smaller for every heavier hole: 4.4×10^{-45} for Cygnus X-1, 1.0×10^{-55} for Sagittarius A* (Computation 154 on the 2.3 nm bound of its day, re-evaluated on $d_0 = 2.2$ am by Computation 216). The expansion first fails where $k\,d_0 = 1$, at

$$r_{\mathrm{break}} \;=\; \left(12^{1/4}\,d_0\right)^{2/3} r_s^{1/3}, \tag{67}$$

Evaluating (67) for a solar mass gives 37 pm, 1.2×10^{-14} of the horizon

radius, enclosing 1.9×10^{-42} of the interior volume, about 10^{25} grains. So the region the projected equations do not cover is a neighbourhood of the central singularity, which item 2 below already lists as open, and not the horizon. There is a floor: $k\,d_0 = 1$ *at* the horizon when r_s falls to $12^{1/4}d_0 \approx 4.1$ am, i.e. $M \approx 2.8 \times 10^9$ kg, and below that mass nothing here applies. The conclusion does not turn on which of PST's two lengths is the cutoff: with the grain ℓ_* in place of d_0 the floor drops to 1.9×10^8 kg, and every astrophysical hole clears both by a wide margin.

What this does and does not buy is worth stating at once. It does not make the horizon identification a theorem; that is a substrate-side claim and is untouched. What it does is remove the *field-equation* objection to treating a black-hole exterior in the projected theory, and that is what §19.9.3 needs: it is why item 1 below is narrower than "no entropy".

8.3.4. What the null threshold fixes about light

The previous subsections used τ_{null} twice, for the light cone and for the horizon. It does more than that, and collecting the uses in one place is worth a page because they are usually presented as four unrelated facts about photons. Computation 160 checks each.

Two polarisations. The physical states of a vector excitation of momentum k are $\{\epsilon : k \cdot \epsilon = 0\}$ up to the gauge shift $\epsilon \sim \epsilon + \lambda k$. The transverse subspace is three-dimensional for every k; the gauge direction can be removed only if k lies in its own transverse subspace, and that is $k \cdot k = 0$ exactly. So a null momentum leaves two states and a timelike one leaves three, confirmed over 200 random momenta of each kind. The counting is ordinary gauge theory. What PST supplies is the input: the photon's momentum is null because the U(1) is unbroken (§13.3), and null is the condition (30) places at $T(\{a, b\}) = \tau_{\text{null}}$.

The equation of state. A photon's worldline is null, so the invariant along it is $s = 0$, and the weight $W(s) = -|s|$ derived at (43) vanishes *identically* on it (§7.4.6, Computation 159). A photon therefore contributes nothing to $\Pi_{\mu\nu}$

from its own worldline, and $w = 1/3$ is forced rather than approached.

The Stefan–Boltzmann law. With $p = \rho/3$ the first law is integrable only if $T \propto \rho^{1/4}$, so $\rho \propto T^4$. That is Boltzmann's argument and it uses no quantum input; PST reaches it because it derives the $w = 1/3$ that the argument needs.

And the coefficient, not merely the scaling. The radiation constant is $a = g\,(\pi^2/30)\,k_B^4/(\hbar^3 c^3)$ with g the number of polarisations. The $g = 2$ that the null condition delivered above gives $7.5657 \times 10^{-16}\,\mathrm{J\,m^{-3}\,K^{-4}}$, the tabulated value to a part in 10^4 (Computation 160); $g = 3$ would miss it by half. So one object fixes the light cone, the state count, the equation of state and the constant in front of T^4.

And the spectrum itself. The four facts above stop short of Planck's law, which this book assembled nowhere. The next subsection assembles it, and its point is less the formula than the accounting: which of the inputs PST forces and which it inherits.

8.3.5. The Planck spectrum, assembled

Writing the derivation down takes five inputs. Four are forced by the structure already in place; the fifth is inherited from ordinary field theory, and saying which is which matters more than the algebra. Computation 161 checks every step.

Forced: the linear dispersion, and the two polarisations. Both are above. Masslessness at τ_{null} gives $\omega = ck$, and the null condition gives $g = 2$.

Forced: the occupancy is bosonic. The photon is an inner fluctuation of the gauge sector and is of even Clifford grade (Computation 108), so it is an even-order composite in the canonical anticommutation relations (CAR) algebra, and even-order composites are bosonic (§11.3). The ladder therefore runs $n = 0, 1, 2, \ldots$ and not $n \in \{0, 1\}$. On the fermionic branch the mean occupation would be $1/(e^x + 1)$ and the spectrum would not be Planck's.

Forced, and the one that is easy to miss: the chemical potential vanishes. §12.1

keeps two discrete structures apart on purpose. The conserved U(1) charge is the *winding* number of $\pi_1(\mathcal{V}_{\hat{\eta}}) = \mathbb{Z}$, the slice $\mathcal{V}_{\hat{\eta}} \cong S^1$ (the 6-sphere $\mathcal{V}$ itself is simply connected); the *occupation* number of the Fock space is an integer unrelated to it. Because the conserved charge is the winding number and not the occupation number, nothing in PST conserves photon number, there is no multiplier to attach to it, and $\mu = 0$ exactly. That separation was made for other reasons, and this is where it does visible work.

Inherited: the ladder itself. That a mode's energy comes in steps of $\hbar\omega$ at all follows from canonical quantisation of the projected Lorentzian action (123), which §12.1 performs in the ordinary way once sublimation has supplied the timelike direction. PST supplies the arena, not the ladder. This book has said elsewhere that the quantisation into countable quanta was missing; that was wrong, and the accurate complaint is the weaker one, that it is inherited rather than derived.

With those, the mean occupation of a mode is

$$\langle n \rangle \;=\; \frac{1}{e^{\hbar\omega/k_B T} - 1}, \tag{68}$$

the density of states per unit volume is $g\,\omega^2/(2\pi^2 c^3)$, and the product of the two is

$$u(\omega, T) \;=\; \frac{g\,\hbar\,\omega^3}{2\pi^2 c^3}\;\frac{1}{e^{\hbar\omega/k_B T} - 1}, \tag{69}$$

which is Planck's law [208], agreeing with the textbook frequency and wavelength forms.

Two checks that this is the same physics as the rest of the subsection. Integrating (69) over all frequencies returns $\rho = g\,(\pi^2/30)\,k_B^4 T^4/(\hbar^3 c^3)$, which is *exactly* the radiation constant of the previous subsection: the spectrum and the Stefan–Boltzmann coefficient are one statement at two resolutions, and the same $g = 2$ is what makes both come out right. Wien's displacement law follows as well, with $b = hc/(4.965114\,k_B) = 2.8977720 \times 10^{-3}\,\mathrm{m\,K}$, the CODATA value.

Where it could fail, and does not. A conserved photon number would put $\mu \neq 0$ into (68) and distort the spectrum away from (69). The Far Infrared Absolute Spectrophotometer (FIRAS) aboard the Cosmic Background Explorer (COBE) measured the microwave background spectrum over the full mission and bounds that distortion at $|\mu| < 9 \times 10^{-5}$ [209], so the structural prediction is confirmed to a part in 10^4 rather than merely asserted. PST does not fit μ; it has nothing to fit it to.

And what the comparison with the standard route actually is. Quantum electrodynamics reaches $\mu = 0$ as well, because the interaction vertex does not conserve photon number. The agreement is expected and is not the point. The point is how much each route has to assume to get there: the standard derivation *posits* that the photon is massless, obtains the two polarisations from that plus gauge invariance, imports bosonic statistics from the spin-statistics theorem with its Lorentz-invariance, locality and positive-energy axioms, and needs the coupling in hand before it can say anything about μ. The assembly above posits none of those four. It takes masslessness from τ_{null}, $g = 2$ from the null condition, bosonic occupancy from Boolean exclusion through CAR parity, and $\mu = 0$ from $\pi_1(\mathcal{V}_{\hat{\eta}})$ alone, before any coupling is written down. Only the ladder is shared. A deeper framework earns its place by reproducing the shallower one's results from fewer inputs, and on that count this passage is four ahead.

The photoelectric effect, which splits in two. Equation (69) is a statement about the equilibrium field, and the photoelectric effect is a statement about matter absorbing a single quantum. That much was said here before, with the conclusion that the effect is absent. It is half absent, and the halves are worth separating, because only one of them needs anything this book does not have.

8.3.6. What the ladder already settles about the photoelectric effect

The kinematics need the ladder and nothing else. If the field's energy comes in steps of $\hbar\omega$, then an electron bound by a work function W that absorbs n of

them leaves with

$$K \;=\; n\,\hbar\omega \;-\; W, \tag{70}$$

and at $n = 1$ that is Einstein's relation [210]. Three consequences follow from energy conservation alone, and Computation 165 checks each: there is a *threshold* at $\omega_0 = W/\hbar$, below which nothing is emitted at any intensity; the slope $\mathrm{d}K/\mathrm{d}\omega$ is $\hbar$ *exactly*; and $\mathrm{d}K/\mathrm{d}I = 0$, so the electron's energy does not depend on how bright the light is.

What makes that a derivation rather than a restatement is the alternative. A classical continuous field delivers energy at a rate $I\sigma$, so it gives $K = I\sigma t - W$: the expression contains no ω at all, hence no threshold, and the energy grows with intensity. Both predictions are wrong, and the second is wrong by a margin that was measured early. Lawrence and Beams bounded the delay between illumination and emission at 3×10^{-9} s in 1928 [211]; for a classical field to deliver a work function to an atomic cross-section that fast it would need an intensity above $10^{10}\,\mathrm{W\,m^{-2}}$, where ordinary photoemission runs near $1\,\mathrm{W\,m^{-2}}$. Ten orders is not a discrepancy to be argued about.

Checked as numbers, not only as a shape. With $hc = 1239.84$ eV nm the threshold wavelength $\lambda_0 = hc/W$ is 579 nm for caesium, 525 nm for sodium, 286 nm for zinc and 219 nm for platinum, each within a nanometre of tabulation (Computation 165).

And the theory says where its own picture stops. One quantum at a time is an intensity regime rather than a law: the Keldysh parameter reaches unity near $10^{14}\,\mathrm{W\,cm^{-2}}$ at 800 nm, which is the measured crossover from multiphoton absorption to tunnelling, and (70) carries the n that describes it. A framework predicting single-photon absorption at every intensity would be wrong, and this one does not.

What is supplied. Of the five ingredients the standard route uses, PST posits none. Four are outright: the ladder, by canonical quantisation on an arena it derives (§12.1); the masslessness that makes the quantum $\hbar\omega$, from τ_{null}; the $\hbar$ in the slope, which is the canonical commutator's, derived from the

Boolean gradient (§10.6); and the electron's fermionic character, from Boolean exclusion through the canonical anticommutation relations (§11.2).

The fifth is the interaction vertex, and it is scoped rather than absent. Its *form* is forced: §13.3.1 shows the local-phase variation of the free Dirac Lagrangian is exactly (156), one term carrying the vector current, so the counterterm and even the gauge transformation law are outputs rather than choices, and the one alternative fails on both the current and its conservation (157). Its *strength* is not supplied: the electromagnetic coupling is not fixed anywhere in this book (§13.3.1).

So the residue was a calculation rather than a missing ingredient, and it is done. Computation 189 takes the vertex of §13.3.1 into the non-relativistic limit, where the vector current couples to the electron's momentum, and applies Fermi's golden rule to a hydrogenic ground state: the photoionisation cross-section is $\sigma = (2^8\pi/3)\,\alpha\,a_0^2\,k^3/(1+k^2a_0^2)^5$ with $\hbar\omega = I_{\mathrm{H}}(1+k^2a_0^2)$, which falls as $(I_{\mathrm{H}}/\hbar\omega)^{7/2}$ at high energy and is the high-energy limit of the exact Coulomb-wave result [215]; the Coulomb final state supplies the threshold value, 6.30 Mb for hydrogen. The axial alternative gives no photoeffect at this order at all, its non-relativistic limit being the pure spin operator $\boldsymbol{\sigma}\cdot\mathbf{A}$, whose matrix element between orthogonal eigenstates vanishes, so the form gauge covariance forces is also the form the effect needs. The rate is σ times the photon flux: at $1\,\mathrm{W\,m^{-2}}$ and 20 eV it is 7×10^{-5} per second per atom with no accumulation, so a square centimetre of cathode emits its first electron within 1.4×10^{-11} s, inside the Lawrence–Beams bound, where the classical field needs 248 s per atom. What the calculation takes as given is α: the coupling's value is not fixed in this book (§13.3.1), and the cross-section inherits that and nothing else.

8.3.7. The Hawking temperature, and four things not claimed

The Hawking temperature of a black hole [46, 47] acquires a precausal reading as the thermal fluctuation of near-threshold configurations at the horizon: configurations lying arbitrarily close to the null-threshold condition fluctuate

across it, and the projection of that fluctuation is outward flux. This is consistent with Jacobson's thermodynamic derivation of the field equations [19], and it grounds that derivation in the projection structure of Π rather than in a phenomenological entropy argument. The operative word is *grounds*: it does not replace it.

Four things are deliberately not claimed.

1. **No microstate count.** The *thermodynamic* Bekenstein–Hawking entropy does follow, given the scope statement above: Wald's Noether-charge construction returns $S = A/4G$ for the Lagrangian (50), and the higher-curvature corrections leave that value exactly unchanged (Computation 154, §19.9.3). What is not derived is the *statistical* entropy. A substrate account of it would have to count near-threshold configurations against horizon area, and PST has no such count. *Nor does the field, in general*: Wald's construction contains no microstate count either, string theory's count runs for extremal and near-extremal holes rather than for the general case, and loop quantum gravity fixes its Barbero–Immirzi parameter *by* matching the entropy it is meant to predict (Table 2). The absence is the field's position, not a deficit peculiar to PST, and what PST adds is that the thermodynamic value survives the higher-curvature corrections exactly.

2. **No result on the central singularity.** The comparison of Table 2 says the *initial* singularity dissolves, because loci are not primitive. That argument concerns the beginning of the instantiated manifold, and it does *not* transfer to a black hole's central singularity without being made again. It now is, in a different form (Computation 216): a point carries at most seven properties (Computation 212), so a collapsed body's content cannot occupy one, and the packing bounds give the minimal diameter of a region holding n properties in local units. That excludes the singularity as a point, and no more: what the region does, the projected equations do not reach, and the theory does not say.

3. **Nothing about the information question.** Whether the outward flux

is unitary is untouched by anything above.

4. **No place in the emergence chain.** Black holes are not a step of the eight-step chain of Chapter 6. They are a limiting case of what step 8 delivers, and the chain is complete without them.

What this section does deliver is the answer to the question a reader asks first, and the reason that answer is forced rather than chosen: PST has exactly one object that can do this work, and it is the same one that made the geometry Lorentzian to begin with.

8.4. Stress-energy as projected tension

The unification of geometry and matter as coprojections of the same precausal tension resolves a longstanding conceptual problem in theoretical physics. In General Relativity, the statement "matter curves spacetime" is taken as a foundational description: matter and spacetime are separate ontological kinds, and one acts on the other. In Quantum Field Theory, fields and their excitations are defined on a background spacetime, which is treated as a fixed classical arena. The two descriptions are mutually inconsistent at a deep level: the quantum fields of quantum field theory (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.

8.5. Newton's constant as a consistency relation

Equation (74) 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 $G = 6\pi d_0^2/(c\,r_0^2)$. Here $\mathcal{N}$ is the trace over the finite internal factor, Λ is the spectral-action cutoff that the observed G fixes (below), not the coherence length d_0, which enters through the kernel normalisation α_G, and $r_0 = v$ is the electroweak vacuum. The full Chamseddine–Connes heat-kernel derivation of the Einstein–Hilbert action and the identification of Newton's G as the Λ^2 moment coefficient is carried through in Computation 5, with the cutoff-function moments f_0, f_2, f_4 evaluated explicitly and the $\mathcal{N}$ multiplier located on the finite internal factor of the spectral triple. 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 spectral cutoff Λ is high and the internal multiplicity $\mathcal{N}$ is finite. The weakness of gravity is therefore the statement that the cutoff sits far above the electroweak scale, not a tuned coincidence. What Λ is in terms of the coherence length is the calibration α_G of $G = \alpha_G d_0^2$, which §8.9 fixes through the kernel normalisation (equation (72)): read with $\alpha_G \sim 1$ the relation would put d_0 near the Planck length, and that reading is not a result of this book, d_0 being the condensate's coherence length rather than its grain (§15).

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 (74) 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 §7.1 with the trace over the finite internal factor. The chain "$(D, \delta) \to$ Boolean Dirichlet form $\to$ 4D Laplace–Beltrami via Mosco $\to$ spectral action $\to G$" runs through the Mosco-limit step (§7.8), which supplies the four-dimensional Lorentzian spin geometry, under A1–A7. What keeps the *value* of G from being parameter-free is the spectral-action calibration, the kernel normalisation and the coherence scale, not the geometry.

8.6. Uniqueness of $G_{\mu\nu}$: 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 Chapter 4.

Step 1: Noether's second theorem. The projected action $S[g]$ inherits diffeomorphism invariance from the derivation of Chapter 4: any two choices of realisation map ρ produce diffeomorphic instantiated geometries, so $S[g]$ is invariant under the infinite-dimensional group $\mathrm{Diff}(M)$. By Noether's second theorem [21], invariance of an action under a local (gauge) symmetry group implies that the Euler–Lagrange expressions satisfy differential identities, the Noether identities. These hold off-shell, as an algebraic consequence of the symmetry, not only on solutions. For diffeomorphism invariance the Noether identity reads:

$$\nabla^\mu \left(\frac{\delta S}{\delta g^{\mu\nu}} \right) = 0 \tag{71}$$

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 (71) 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 satisfy three conditions.

(i) It is symmetric, since it arises by varying a diffeomorphism-invariant scalar action with respect to the symmetric tensor $g^{\mu\nu}$.

(ii) It is divergence-free, by the Noether identity (71).

(iii) It is 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 (49) 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 (50) 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 (49) 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.

8.7. The equivalence principle from single-source projection

General Relativity postulates the weak equivalence principle: that all matter falls identically in a gravitational field, regardless of composition; this enters 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 $\nabla^\mu T_{\mu\nu} = 0$ together with the standard pole-dipole test-body limit, derived from the Noether identity (71). 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.

8.8. The gravitational constant as a unit conversion

The factor $8\pi G$ in equation (48) 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 quantise 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.

8.9. G from the projection scale: form and magnitude

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 [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 (31) 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 realisation map ρ carries a characteristic length: the smallest separation $d_0 > 0$ at which the projection produces a smooth geometry well-approximated by a Lorentzian manifold, below which the continuum description of Π ceases to hold. As Chapter 15 makes precise, d_0 is not the microscopic grain of the substrate but the coherence length of the modal condensate, a distinct scale that near a second-order transition sits many orders of magnitude above

the grain; it is the finest resolution at which the instantiated geometry can smoothly 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 (31) is evaluated, the measure $\mu(C')$, dimensionless as derived in §2.5 via the canonical Bernoulli measure of equation (9), 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 (48) yields a dimensionless coefficient α_G fixed by the detailed form of K:

$$G \;=\; \alpha_G \cdot d_0^2 \qquad (\hbar = c = 1) \tag{72}$$

Since $G = \ell_P^2$ in these units, equation (72) 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. This is the same relation as the spectral-action closed form $G = 6\pi d_0^2/(c\, r_0^2)$ of equation (74) reparametrised, with $\alpha_G = 6\pi/(c\, r_0^2)$; the r_0- and c-dependence there is exactly the kernel content that fixes the otherwise dimensionless α_G. Chapter 15 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, 2.2 am under the canonical normalisation, whose Casimir correction is of derived form but far below present reach. 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 (72) 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 $M_* = 4\pi m_h \sqrt{2/3}$ 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.

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

Equation (61) 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,

$$\mathcal{N} \;=\; \mathrm{Tr}_{\mathcal{H}_F}(\mathbf{1}), \tag{73}$$

the fermionic multiplicity of the internal Hilbert space (one generation's or the full fermion count; the choice is absorbed into c below). 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 (61) reads

$$G \;=\; \frac{6\pi\, d_0^2}{c\, r_0^2}. \tag{74}$$

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. (The quartic $b = \lambda_{\mathrm{PST}} = 1/4$ of P3 is stated for the dimensionless order parameter; with r_0 in physical units, $r_0 = v$ and $m_h^2 = 4\varepsilon$ fix $b = m_h^2/(8v^2) = 0.032$, the two related by the canonical normalisation of §12.2, Computation 197.)

Dimensional note. In natural units ($\hbar = c_{\mathrm{light}} = 1$), 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 §8.5.

8.11. Summary

This chapter checks that the projected geometry is Einstein's. Curvature is the non-uniformity of the tension distribution (§8.1); the weak-field limit recovers Poisson's equation and Newtonian gravity (§8.2); in the strong-field limit a horizon is a threshold surface and light cannot leave it because the null threshold is a sign cut, with the Hawking temperature given a precausal reading and four things explicitly not claimed (§8.3). The Einstein tensor is the unique choice by diffeomorphism invariance and Lovelock's theorem (§8.6); the equivalence principle is a consequence of single-source projection rather than a postulate (§8.7). Newton's constant is not a property of the substrate but of the projection: a unit conversion between substrate and laboratory lengths, $G = \alpha_G d_0^2$, with its magnitude fixed through the kernel normalisation $\mathcal{N}$, the trace over the finite internal factor, as a consistency relation among (m_h, M_P, v) (§8.5, §8.10).

9. Vacuum Structure and Angular Momentum

Take everything out of a box and something is still left. This chapter works out what the substrate leaves behind when nothing is there, and finds it has structure: enough to explain why things are able to spin, and enough to matter for how the universe expands.

9.1. The cosmological constant: substrate vs emergent energy

The Landau–Ginzburg valley energy of the modal vacuum is

$$\mathcal{F}\big|_{\mathcal{V}} = \int_{\mathcal{C}} \left[-\frac{(|T(C)| - \tau)^2}{4b} \right] d\mu(C) \;<\; 0, \tag{75}$$

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 Landau–Ginzburg (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 $\mathcal{F}\big|_{\mathcal{V}} < 0$ and $\Lambda_{\text{obs}} > 0$ dissolves once that distinction is made precise. This section addresses the sign of Λ through two complementary mechanisms: A, a calculable bare zero-point value whose sign is regulator-dependent, and B, a structural $w = -1$ contribution that carries the sign conditional on the substrate's pre-spatial framing. 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.

9.1.1. What $\mathcal{F}\big|_{\mathcal{V}}$ is, and is not

The static valley depth $\mathcal{F}\big|_{\mathcal{V}}$ 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 §16.2); the projection Π defines what couples to $g_{\mu\nu}$, not the bare LG potential. $\mathcal{F}\big|_{\mathcal{V}} < 0$ is therefore a correct statement about the LG vacuum but not, on its own, a statement about Λ.

9.1.2. Mechanism A: bare Goldstone zero-point on the slice $\mathcal{V}_{\hat{\eta}} \cong S^1$

The Goldstone phase θ of the broken U(1) on $\mathcal{V}$ is a massless field whose bare zero-point fluctuations at a fixed ultraviolet (UV) regulator Λ_{UV} contribute

$$\rho_{\mathrm{A}}^{\mathrm{bare}} \;=\; \frac{1}{16\pi^2}\,\Lambda_{\mathrm{UV}}^4 \;+\; \text{counter-terms}, \tag{76}$$

which is positive in absolute value at any naive cutoff and dominates $|\mathcal{F}\big|_{\mathcal{V}}|$. At PST's modal cutoff $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV the bare contribution, taken at the un-suppressed naive-cutoff scale $\Lambda_{\mathrm{UV}} \sim M_*$ (dropping the loop factor of equation (76) for an order-of-magnitude estimate), is $\rho_{\mathrm{A}}^{\mathrm{bare}} \sim M_*^4 \approx 10^{12}$ GeV^4, larger than $|\mathcal{F}\big|_{\mathcal{V}}| \sim \varepsilon^2 \sim 10^8$ GeV^4 by four orders, and larger than the renormalised observed value $\Lambda_{\mathrm{obs}}/(8\pi G) \approx 10^{-47}$ GeV^4 by some 60 orders. Matching Λ_{obs} requires the standard quantum field theory (QFT) with-gravity counter-term cancellation between $\rho_{\mathrm{A}}^{\mathrm{bare}}$, $\mathcal{F}\big|_{\mathcal{V}}$, and the bare cosmological constant. This is the standard cosmological-constant fine-tuning problem of any quantum field theory (QFT) coupled to gravity; PST inherits it through the LG–Higgs identification, and does not claim to resolve it.

Scope of Mechanism A. It might be supposed that Mechanism A *delivers the sign* of Λ_{obs}. That would be too strong: the sign of the bare zero-point energy is in fact regulator-dependent (this is the heart of the cosmological-constant problem), and once a large counter-term cancellation is conceded for the magnitude, the sign of Λ_{obs} is set by the counter-term, not by $\rho_{\mathrm{A}}^{\mathrm{bare}}$. Mechanism A therefore contributes a *calculable* bare value at fixed regulator (positive in the naive $\Lambda_{\mathrm{UV}}^4/(16\pi^2)$ form), but does not on its own pin the sign of the renormalised observed Λ. The sign of Λ_{obs} comes from Mechanism B

below, subject to its own caveats.

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 $J^0 = 2r_0^2\omega$ is not the electromagnetic charge, which annihilates H^0 and does not move θ (§8.2.1); it is the charge of the broken neutral direction the Z absorbs. In unitary gauge the rotation is a constant Z^0 background, a preferred-frame weak field with energy density $r_0^2\omega^2$ that nothing observes and that is not the energy minimum: the vacuum has $\omega = 0$ because it is the vacuum, and $J^0 = 0$ with it.

9.1.3. Mechanism B: the ongoing-sublimation source

A structurally independent contribution arises from the cosmological continuity equation derived in §16.5. The substrate maintains an equilibrium density j_τ of threshold-crossed configurations in its pre-geometric layer, a substrate-internal density measured against the configuration-space measure $\mu = \bigotimes_{a\in D} \mathrm{Bern}(\frac{1}{2})$ of P3, not against any emergent-spacetime proper volume (the substrate has no spatial extent against which a proper-volume rate could be defined). Each threshold crossing carries the energy quantum $\Delta\mathcal{F} = \mathcal{F}_{\mathrm{pre}} - \mathcal{F}\big|_{\mathcal{V}} > 0$ (the configuration descends from the high-tension pre-threshold state to the valley). The modal-sublimation functor $\Phi : \mathbf{S} \to \mathbf{G}$ carries this substrate-internal saturation into emergent spacetime as a true energy density $\rho_{\mathrm{inst}} = j_\tau\, \Delta\mathcal{F}/c^2 > 0$ that neither dilutes nor redshifts, so the covariant continuity equation forces $w = -1$. This derivation (non-dilution from the substrate's spacelessness, the continuity equation, and the positive sign from $\Delta\mathcal{F} > 0$) is carried out in full in §16.5, its canonical home; here we note only what is distinctive to Mechanism B. The choice that makes the equation of state $w = -1$ rather than $w = 0$ is the definition of j_τ as a saturation density in the pre-spatial substrate (measured against the configuration-space measure of P3, with no proper-volume rate to refer to) rather than as a per-emergent-volume creation rate: the latter would give Hoyle–Bondi–Gold-style steady-state ($w = 0$, continuous matter creation compensating dilution), not a cosmological constant. This is forced by the pre-spatial framing of the substrate, and is

what separates PST's mechanism from the classic steady-state models.

Scope of Mechanism B. One might present the no-dilution property as a derived structural consequence of the substrate's spacelessness. This is in fact close to circular: the assertion "the substrate has no spacetime extent, so its outputs do not dilute under expansion" is itself the substantive content of the pre-spatial framing of $\mathcal{S}$, not a derived corollary of it. Mechanism B is therefore best read as: *conditional* on the pre-spatial framing of the substrate (an integral part of P1–P3 and the modal-sublimation functor $\Phi : \mathbf{S} \to \mathbf{G}$, but not derived from any narrower fact), the ongoing-sublimation contribution to the energy budget of emergent spacetime carries $w = -1$ exactly. This places Mechanism B in the same family as the Hoyle–Narlikar steady-state C-field [147] and similar ongoing-creation proposals; that family is constrained by observation (cosmic microwave background [CMB] blackbody, nucleosynthesis), as discussed at the cosmology level in §16.10. The PST formulation differs by treating ongoing instantiation as a structural feature of the substrate-to-spacetime projection rather than a phenomenological addition to general relativity (GR).

The two mechanisms differ both on whether they deliver the sign and on which is dominant:

- Mechanism A contributes a calculable bare zero-point energy at fixed regulator. Its sign is regulator-dependent, and the observed sign of Λ_{obs} is set by the QFT counter-term, not by the bare Mechanism A value. PST inherits the Standard Model (SM) cosmological-constant magnitude problem through this mechanism.
- Mechanism B contributes $w = -1$ structurally, with the positive sign of ρ_{inst} following from the direction of threshold crossing. The non-dilution premise is the pre-spatial framing of the substrate, not an independent derivation; conditional on that framing, the equation of state is fixed.

The honest joint statement is therefore: PST does *not* derive Λ_{obs} from first principles. What it offers is (i) a calculable bare zero-point contribution (Mechanism A) whose sign is renormalisation-scheme-dependent, and (ii) a

structural $w = -1$ equation-of-state contribution (Mechanism B) conditional on the substrate's pre-spatial framing. Whether Mechanism B is the dominant contributor to Λ_{obs} is an open question; if it is, the small observed value would not require the SM fine-tuning, but a structural argument for the absolute magnitude $j_\tau\,\Delta\mathcal{F}$ would still be needed. GR's metric and curvature emerge from the spectral action (Computation 5), while the cosmological-constant *magnitude* is renormalised just as in any QFT coupled to gravity.

One route to that magnitude from the theory's own resources was asked about and is closed (Computation 188). The second edition derives no area law for Π (Table 2, §19.9.3), and one written with the area quantum set by a length the theory owns would be false rather than weak: the energy bound it implies, that a region holds no more vacuum energy than a black hole of its size, is 10^{-32} of the critical density with the grain ℓ_* and 5×10^{-35} with the derived d_0 (5×10^{-53} at the 2.3 nm bound Computation 188 ran at), against an observed ρ_Λ of 0.7 of it. The bound that lands is general relativity's, reached through the emergent G of (61); it is an inequality on the total, and since $\Delta\mathcal{F}$ enters only through the product $j_\tau\,\Delta\mathcal{F}$ it fixes neither factor. The finiteness of the internal structure sets a scale, M_*^4, fifty-nine orders above ρ_Λ, which is Mechanism A again. Nothing in the theory's own equipment reaches $\Delta\mathcal{F}$, and nothing here needs a fourth postulate or an imported bound.

9.2. The degenerate vacuum manifold and angular momentum

The order parameter ψ is V_7-valued (P3, §1.7). For the angular-momentum analysis of this section 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:

$$\mathcal{F}(\psi_{\mathbb{C}}, C) = \int_{\mathcal{C}} \Big[-\varepsilon \, |\psi_{\mathbb{C}}|^2 + b \, |\psi_{\mathbb{C}}|^4 + c \, |\nabla_{\mathcal{C}} \psi_{\mathbb{C}}|^2 \Big] d\mu(C) \tag{77}$$

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:

$$\mathcal{V}_{\hat{\eta}} = \big\{ \, \psi_{\mathbb{C}} \in \mathbb{C}_{\hat{\tau},\hat{\eta}} \; : \; |\psi_{\mathbb{C}}| = r_0 \, \big\}, \qquad r_0 = \sqrt{\frac{|T(C)| - \tau}{2b}} = \sqrt{\frac{\varepsilon}{2b}} \tag{78}$$

The slice $\mathcal{V}_{\hat{\eta}}$ of the vacuum manifold is topologically S^1: the full $\mathcal{V} \cong S^6 \subset V_7$ of §9.1 meets the complex slice $\mathbb{C}_{\hat{\tau},\hat{\eta}}$ in a circle. 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 precausal excitation manifests as a propagating mode in the instantiated geometry.

One qualification, and the photon does not depend on it. The degeneracy described in this paragraph is the U(1) slice of a picture in which the order parameter's direction in V_7 is *determined* by P2 rather than spontaneously chosen. Those angular directions therefore supply no Goldstone, and the

symmetry that survives fixes the vacuum and so supplies none either (Computation 170). The mode described here is the *phase* of the complexified order parameter, which determination cannot touch because $\hat{\tau}$ is a real direction and carries no phase (Computation 171).

And it is not the photon, which is worth saying plainly because the sentence above reads as though it were. Both alternatives fail: left as a free massless *scalar* below M_* it would break F4 on its own, the Standard Model having no massless scalar, and a U(1) that is both gauged and broken gives a *massive* boson, which is a superconductor rather than electromagnetism. The consistent reading is the Standard Model's own: θ is *eaten* in the electroweak Higgs mechanism, among the would-be Goldstones taken up by $W^{\pm}$ and Z, and the photon is the *unbroken* combination, massless because that generator is unbroken. So θ is the mode whose wave equation §4.12 argues from, described before it is eaten, and the photon is what the gauging step below delivers.

Which component it is, concretely. With $\psi = H^0$ the substrate's complex line $\mathbb{C}\hat{\tau}$ carries two real degrees of freedom and they are the neutral component's: the modulus is the Higgs and the phase is the would-be Goldstone the Z takes up (Computation 172). The doublet's other two, $H^{\pm}$, come from $\mathrm{SU}(2)_L$ rather than from the complex line, which is consistent, since $\hat{\tau}$ is a direction and a phase and carries no charge. Two and two make the doublet's four.

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. That step is carried out constructively in §13.3: the projection Π maps the continuum of substrate configurations to the full manifold M, so independent configurations at different loci of M select independent vacuum phases $\theta(x)$; the symmetry parameter is therefore an arbitrary smooth function $\alpha(x)$ and the symmetry is local, with $\theta(x)$ a section of a U(1) principal bundle and $D_\mu\theta = \partial_\mu\theta - A_\mu$ its

covariant derivative. *What that step delivers is worth stating exactly, since an earlier version of this paragraph had it backwards.* The connection it introduces couples to a phase whose U(1) is *broken*, $|\psi| = r_0 \neq 0$, and a gauged broken U(1) gives a *massive* vector: in unitary gauge θ is eaten and the Stueckelberg term $r_0^2 A_\mu A^\mu$ is what remains, a superconductor rather than electromagnetism. Once the electroweak embedding of §13.4 is in place that vector is the Z (Computation 172). Nor is the Maxwell term (155) obtained by integrating out θ; it is supplied by the inner fluctuations of the spectral action (Computation 6). The photon is the connection of the *unbroken* combination $Q = T_3 + Y/2$, massless because Q annihilates the vacuum, which is the Standard Model's own argument and the one this book inherits (§13.3); masslessness does not follow from the inhomogeneous transformation law of A_μ on its own, since gauge invariance permits exactly the Stueckelberg mass above. So the route from the phase is constructive and it ends at the Z, not at the photon: θ *is not the photon* (Computation 171, Computation 172), and the photon has one route, as the unbroken connection, with minimal coupling forced by gauge covariance (Computation 166).

9.3. 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 the slice $\mathcal{V}_{\hat{\eta}} \cong S^1$, 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 rotation of galaxies, exhibits stable circular structure. PST requires no additional postulate.

9.4. 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:

$$j^\mu = 2r^2\, \partial^\mu\theta \tag{79}$$

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

$$L = \int j^0\, d^3x = 2r_0^2\, \dot{\theta} \tag{80}$$

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. 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 carried out in the next subsection (§9.5), where the identification $L_z = mr^2\omega$ is obtained as a structural consequence of the Poincaré invariance of the projected Goldstone action.

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, as derived in the next subsection (§9.5).

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, in its integer, the precausal statement that the phase of the order parameter can only wind an integer number of times around the ring without discontinuity; the unit $\hbar$ is supplied by canonical quantisation, not by the winding (below). Spin is the near-threshold manifestation of the same U(1) charge that produces orbital angular momentum in the far-threshold classical regime.

9.5. Angular momentum and the Goldstone–geodesic connection

The previous subsection established the Noether charge $L = 2r_0^2\dot{\theta}$ and announced that the explicit computation of its projection to physical orbital angular momentum is carried out here. This section delivers it 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 $\theta : M \to \mathbb{R}/(2\pi\mathbb{Z})$ 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

$$T^\theta_{\mu\nu} \;=\; 2r_0^2\,\partial_\mu\theta\,\partial_\nu\theta \;-\; g_{\mu\nu}\,r_0^2\,\partial_\lambda\theta\,\partial^\lambda\theta, \tag{81}$$

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

$$L_{ij} = \int \left(x_i T^\theta_{j0} - x_j T^\theta_{i0}\right) \mathrm{d}^3x = 2r_0^2 \int \left(x_i\,\partial_j\theta - x_j\,\partial_i\theta\right)\dot{\theta}\,\mathrm{d}^3x. \tag{82}$$

These are conserved by the Bianchi identity [10] $\nabla^\mu T^\theta_{\mu\nu} = 0$ (equation (71)), which follows from the diffeomorphism invariance derived in Chapter 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

$$L_z^{\text{shell}} \;=\; 2r_0^2 \int_{r_1}^{r_2} n\,\omega\, 4\pi r^2 \,\mathrm{d}r \;=\; \frac{8\pi}{3}\, n\, r_0^2\, \omega\, (r_2^3 - r_1^3), \tag{83}$$

the integrand being $2r_0^2 n\dot{\theta}$ because $x\partial_y - y\partial_x$ is ∂_φ and $\partial_\varphi(n\varphi) = n$ is dimensionless. This scales as r^3, so the far-field Goldstone contribution is not localised: it grows without bound with the outer radius. *That divergence is not removed by the short-distance cutoff*, being infrared rather than ultraviolet, and it is the classic global-vortex behaviour. The winding label is an integer,

$$n \;\in\; \mathbb{Z}, \tag{84}$$

because $\pi_1(\mathcal{V}_{\hat{\eta}}) = \pi_1(S^1) = \mathbb{Z}$ on the slice. *The topology fixes the integer and not the unit.* For $\theta = n\varphi$ a spatial rotation by α acts as an internal phase shift by $n\alpha$, so $L_z = nQ$ with Q the internal charge, and single-valuedness under a 2π rotation then gives $L_z \in \hbar\mathbb{Z}$. But that last step is canonical quantisation, and $\hbar$ enters through equation (106) of §10.6, where it is supplied as the unit of momentum rather than derived. *PST supplies the vacuum topology that makes the spectrum discrete; it does not supply $\hbar$* (Computation 161).

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 $L_z = n\hbar$, the winding integer of equation (84) times the canonical unit, in the classical ($\hbar \to 0$, $n \to \infty$) limit requires

$$m \;=\; \lim_{n\to\infty} \frac{n\hbar}{r^2\omega} \;=\; \frac{2r_0^2}{r^2} \int \dot{\theta}\, \mathrm{d}^3x_\perp \tag{85}$$

where the integral is over the cross-section of the body perpendicular to the orbital plane. This is consistent with the stress-energy identification $m = \int T_{00}\, \mathrm{d}^3x$ 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

$$S_\theta \;=\; r_0^2 \int_M \partial_\mu\theta\, \partial^\mu\theta\, \sqrt{-g}\, \mathrm{d}^4x \tag{86}$$

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 $P^\mu = \int T^{0\mu}\, \mathrm{d}^3x$ is a Lorentz 4-vector [97]. A global Goldstone vortex is *not* a finite-energy configuration: with $|\nabla\theta| = n/\rho$ the energy per unit length is $2\pi r_0^2 n^2 \ln(R/d_0)$, logarithmically divergent in the outer scale R, and d_0 regulates only the core. The soliton momentum theorem needs finite total energy, so it applies to the *regulated* configuration, with R set by the inter-defect spacing, and the mass it delivers is correspondingly cutoff-dependent. With that understood, and only with it [97]: $\mathbf{P} = \gamma m\mathbf{v}$, where $m = E_{\text{rest}}/c^2$ is the rest mass. In the non-relativistic limit $v = r\omega \ll c$:

$$\mathbf{P} \;=\; m\mathbf{v}, \qquad v = r\omega. \tag{87}$$

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

$$L_z \;=\; r\,P_\phi \;=\; mr^2\omega. \tag{88}$$

Including the intrinsic topological contribution $n\hbar$, the winding integer of equation (84) times the canonical unit:

$$L_z^{\text{total}} \;=\; mr^2\omega \;+\; n\hbar. \tag{89}$$

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 $T^\theta_{00} = r_0^2|\nabla\theta|^2$, so

$$mc^2 \;=\; r_0^2\int_M |\nabla\theta|^2\,\mathrm{d}^3x. \tag{90}$$

For a localised phase lump of width d_0 with linear profile $\theta(\rho) = \pi\big(1-\rho/d_0\big)$ for $\rho \le d_0$ (and $\theta = 0$ beyond), the gradient is uniform inside the core: $|\nabla\theta| = \pi/d_0$. The integral evaluates to

$$\int_{|\mathbf{x}|\le d_0} |\nabla\theta|^2\,\mathrm{d}^3x \;=\; \frac{\pi^2}{d_0^2}\cdot\frac{4\pi d_0^3}{3} \;=\; \frac{4\pi^3 d_0}{3}, \tag{91}$$

giving the PST elementary mass scale

$$m_0 \;=\; \frac{4\pi^3}{3}\cdot\frac{r_0^2\,d_0}{c^2}. \tag{92}$$

What this is, and what it is not. It is a dimensional scale: any phase con-

figuration localised on a width d_0 with an $O(1)$ excursion has rest energy $\propto r_0^2 d_0$, and $m_0 \propto r_0^2 d_0/c^2$ is robust. It is *not* the mass of a vortex quantum, and the profile above is not a vortex: $\theta(\rho)$ depends on the radius alone, so it is constant on every closed loop and its winding number is zero. Nor could it be, since $\pi_2(S^1) = 0$ and a U(1) order parameter admits no point or spherical defect in three dimensions; the stable defects are lines, classified by $\pi_1(S^1) = \mathbb{Z}$, and those are exactly the configurations whose energy diverges above. The coefficient $4\pi^3/3 \approx 41.3$ is exact for this profile and carries no further significance. *The orbital angular-momentum treatment is therefore closed on the terms it can be*: it supplies the tensor, the integer winding label and a mass scale, $L_z = nQ$ is an identity for every axisymmetric configuration with Q the Noether charge, and the unit is the absolute normalisation of the Goldstone stiffness r_0^2, which is the gradient-term normalisation Appendix C carries as the canonical-normalisation premise (Computations 194, 201); $\hbar$ is inherited at (106), and the treatment is carried there rather than as an item of its own.

9.6. Summary

This chapter works out what the substrate leaves behind when nothing is there. The modal vacuum has negative Landau–Ginzburg valley energy, which the chapter keeps apart from the emergent energy the instantiated sector sees; the sign of the cosmological constant comes from the ongoing-sublimation term, conditional on the pre-spatial framing (§9.1), and its magnitude is configuration content, not structure. The vacuum manifold is the degenerate 6-sphere $\mathcal{V} \cong S^6 \subset V_7$, read in the 2-plane slice spanned by $\hat{\tau}$ and $\hat{\eta}$ as the ring of minima (§9.2), which standard physics, with its unique symmetric ground state, lacks. The global phase symmetry of the Goldstone mode survives projection, and its Noether charge $L = 2r_0^2\dot{\theta}$ is orbital angular momentum (§9.4), delivered explicitly through the Goldstone action and the geodesic condition (§9.5); its unit is inherited at (106) and is carried with the normalisation premise.

10. Recovering Quantum Mechanics

Quantum mechanics is the best tested theory in physics, so any deeper theory has to reproduce it exactly. This chapter is PST paying that bill: where the famous probability rule comes from, what actually happens during a measurement, and what the theory says about the experiments every physics student meets, from the double slit to Bell's.

10.1. 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 chapter concerns 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 chapter are therefore predictions about the domain where existing theories break down, approach their limits, or require assumptions that PST renders derivable.

This chapter collects the structures of standard quantum mechanics that PST *recovers* rather than predicts anew: the Born rule, the measurement/collapse rule, and the canonical commutation relations are derived from the substrate

machinery, confirming that PST reproduces known quantum mechanics rather than contradicting it. The genuinely distinguishing prediction, the short-distance d^{-6} Casimir correction through which PST departs from standard physics and can be tested at laboratory scale (the collider-side tests, the Higgs quartic at M_* and the Higgs-p_T deficit, are in Chapter 19), has a chapter of its own (Chapter 15).

10.2. The Born rule from the Bernoulli measure

A foundational question the projected field theory (§12.1) leaves to this chapter 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 Chapter 4.

The PST Hilbert space. The gradient structure derivation of Chapter 4 established that the space $L^2(\mathcal{P}(D), \mu)$ of square-integrable complex functionals on configuration space, with the Bernoulli measure $\mu = \bigotimes_a \mathrm{Bern}(1/2)$, is a separable Hilbert space. The inner product is

$$\langle \phi, \psi \rangle \;=\; \int_{\mathcal{P}(D)} \phi^*(C)\, \psi(C)\, \mathrm{d}\mu(C), \tag{93}$$

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 $P_i : L^2 \to \mathcal{H}_i$ 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 Chapter 4): 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:

$$P(\text{outcome } o_i \mid \psi) \;=\; \frac{\displaystyle\int_{C\in\mathcal{S}_i} |\psi(C)|^2 \,\mathrm{d}\mu(C)}{\displaystyle\int_{\mathcal{P}(D)} |\psi(C)|^2 \,\mathrm{d}\mu(C)} \;=\; \frac{\|P_i\psi\|^2}{\|\psi\|^2} \;=\; \|P_i\psi\|^2 \qquad (94)$$

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: $\int_{S_i\cup S_j} |\psi|^2 \,\mathrm{d}\mu = \int_{S_i} |\psi|^2 \,\mathrm{d}\mu + \int_{S_j} |\psi|^2 \,\mathrm{d}\mu$ for disjoint S_i, S_j. By Gleason's theorem [54] applied to the separable Hilbert space $L^2(\mathcal{P}(D), \mu)$ (dimension ≥ 3 for $|D| \geq 2$), *every* non-contextual probability measure on the orthogonal projectors has the form $P \mapsto \mathrm{Tr}(\rho P)$ for some density operator ρ (Computation 135). That is a real constraint and not a vacuous one: it excludes every non-quantum probability assignment. *But it does not by itself single out* $P \mapsto \|P\psi\|^2$, and the difference matters here. The threshold-crossing mechanism above fixes the measure on the *configuration* algebra, that is on one maximal abelian resolution; and a density operator is not determined by its diagonal. For $|D| = 2$ the pure state $\psi = \frac{1}{2}(1,1,1,1)$ and the maximally mixed $\rho = \frac{1}{4}\mathbb{1}$ assign identical probabilities to every configuration and differ by a factor of four on the projector onto ψ itself.

What closes the gap is the pointer, not Gleason. A measurement of an arbitrary $\hat{O}$ terminates in a reading of an apparatus pointer, and the von Neumann coupling (96) carries $\psi \otimes |0\rangle \mapsto \sum_i c_i |o_i\rangle \otimes |\lambda o_i\rangle$, after which the pointer

position *is* a configuration-diagonal quantity. Applying the configuration-basis rule to the coupled state returns $P(o_i) = |c_i|^2 = \|P_i\psi\|^2$ for arbitrary $\hat{O}$. So the Born rule is not an additional postulate of PST: it follows from the Bernoulli measure *together with* the von Neumann coupling and the unitary propagation between crossings, which §10.4 states plainly that PST assumes rather than derives. Gleason constrains the form; it does not supply the state.

10.2.1. Two non-contextualities, and which one is assumed

The appeal to Gleason invites an objection worth answering directly, because the substrate is contextual by construction: P1 makes distinction the primitive, tension is defined by complement-asymmetry, and no elementary property carries a context-free value. How can a derivation that assumes *non-contextuality* sit inside an ontology of that kind?

The two conditions are distinct. Gleason's hypothesis constrains the probability *measure*: the value assigned to a projector P must not depend on which orthogonal resolution of the identity P is taken to belong to, which is frame-function additivity in an arbitrary orthonormal basis (Computation 135). Within PST this is not an additional assumption but the measure-theoretic face of the substrate's label-indifference, the same invariance that fixes μ uniquely (§2.5): a substrate with no favoured property cannot weight a projector differently according to the company it keeps. Kochen–Specker contextuality [193] constrains *value assignments*, that is, whether every observable can be given a definite value independently of what is measured alongside it. These are not the same condition, and they do not merely coexist. For $\dim \mathcal{H} \geq 3$ Gleason's theorem admits no dispersion-free frame function and therefore *entails* that no non-contextual value assignment exists; standard quantum mechanics is the witness that the Born rule is Gleason-non-contextual while the theory is Kochen–Specker contextual.

The accounting is therefore favourable rather than awkward. What is assumed at the Gleason step is even-handedness of the measure, itself already forced by relabelling invariance; what is *derived* is the contextuality of values, which is

the feature the substrate carries by definition. The apparatus-dependence of the modal threshold in §10.3, $\tau \mapsto \tau_i$, is the mechanism by which that derived contextuality is realised measurement by measurement.

The derivation above establishes the Born rule for measurements that decompose into orthogonal eigenspaces in $L^2(\mathcal{P}(D), \mu)$. The full measurement postulate (collapse of the wave function) follows from the threshold-crossing mechanism of P3 through the three closures and one condition of §10.3 (explicit von Neumann measurement [24] Hamiltonian, decoherence inheritance, threshold-crossing uniqueness + irreversibility theorem, and agreement with the Born rule). The fermionic-sector treatment at full functional-analytic rigour follows from the three closures of §11.3, two of them given A1 to A4 (convergence of the canonical anticommutation relations [CAR] to the continuum, Wightman axioms, substrate-derived Dirac equation). The canonical commutation relations hold in §10.6, given the identification of the projected Boolean gradient with ∂_μ.

10.3. The full measurement postulate from threshold-crossing

The Born rule of §10.2 delivers *probabilities*: the probability of measurement outcome o_i given pre-measurement state ψ is $\|P_i\psi\|^2$. The full measurement postulate adds a *state-update rule*: after observing outcome o_i, the post-measurement state is the projection $P_i\psi/\|P_i\psi\|$ ("collapse"). In standard quantum mechanics, both the probability rule and the state-update rule are postulated separately; in PST, the state-update rule emerges from the same substrate-threshold-crossing mechanism that delivers the Born rule.

10.3.1. The threshold-crossing event as collapse

In the substrate-projection picture, a quantum measurement is the moment at which the substrate configuration C crosses the modal threshold for an observable. Concretely, let $\hat{O}$ have spectral decomposition $\hat{O} = \sum_i o_i P_i$ on $\mathcal{H} = L^2(M, d^4x)$ and let $\mathcal{H}_i$ be the eigenspace of o_i. The substrate-projection

of $\mathcal{H}_i$ under Π corresponds to a subset $\mathcal{S}_i \subseteq \mathcal{P}(D)$ of substrate configurations whose projected wavefunction lies in $\mathcal{H}_i$, and for a configuration-diagonal $\hat{O}$ the collection $\{\mathcal{S}_i\}_i$ partitions $\mathcal{P}(D)$. *That restriction is not decorative*: a family of subsets induces multiplication-by-indicator projectors, so no two non-commuting observables can both partition $\mathcal{P}(D)$, which is why §10.4.3 finds the μ_x and μ_z partitions incompatible. For a general $\hat{O}$ the reduction to a partition runs through the pointer, as above.

When a measurement is performed, the substrate's tension structure couples to the measuring device, modifying the modal threshold τ into an outcome-dependent threshold τ_i. The substrate configuration C crosses τ_i for some specific i, not for all i simultaneously, because the threshold-crossing event of P3 is by construction a single irreversible directed transition (the parity-violating mechanism of Computations 8 and 11 underwrites the irreversibility). The selection of i at threshold-crossing is therefore unique, and the post-measurement substrate is in a configuration $C' \in \mathcal{S}_i$ projecting to a wavefunction in $\mathcal{H}_i$: this *is* the collapse.

10.3.2. Probability of selection

The probability that the threshold-crossing event selects outcome o_i is, by the Born rule of §10.2,

$$\begin{aligned}\mathbb{P}(\text{crossing for } o_i \mid \psi) &= \mu\big(\mathcal{S}_i \cap \{C : |\psi(C)|^2 \text{ above threshold}\}\big)/\|\psi\|^2 \\ &= \|P_i\psi\|^2/\|\psi\|^2,\end{aligned} \tag{95}$$

which is the Born probability. The collapse and the Born probability are therefore not two separate postulates but two facets of the same substrate-threshold-crossing event: the event selects outcome o_i with probability $\|P_i\psi\|^2$, and the post-event state is in $\mathcal{H}_i$ by construction.

The measurement postulate, both probability rule and state-update rule, is delivered structurally by the combination of (i) the configuration-basis rule on the Bernoulli measure, Gleason's theorem fixing its form and the pointer

coupling supplying the state (§10.2, giving the probability rule) and (ii) the threshold-crossing irreversibility of P3 (giving the state-update rule). No additional postulate is required beyond the unitary propagation between crossings, which PST assumes (§10.4).

10.3.2.1. Closure (1): Explicit measurement-coupling Hamiltonian. The standard von Neumann measurement model couples system and apparatus via the impulsive Hamiltonian

$$H_{\text{int}}(t) \;=\; \lambda\, g(t)\, \hat{O} \otimes \hat{P}_{\text{app}}, \tag{96}$$

where $\hat{O} = \sum_i o_i P_i$ is the system observable, $\hat{P}_{\text{app}}$ is the apparatus pointer momentum, $g(t)$ is a short pulse of unit integral (idealised as $g(t) = \delta(t)$), and $\lambda \gg 1$ a strong coupling. Under $U = \exp(-i \int H_{\text{int}}\, dt/\hbar) = \exp(-i\lambda \hat{O} \otimes \hat{P}_{\text{app}}/\hbar)$, the joint evolution acts as

$$|\psi\rangle \otimes |0\rangle_{\text{app}} \;=\; \Big(\sum_i c_i |o_i\rangle\Big) \otimes |0\rangle_{\text{app}} \;\xrightarrow{U}\; \sum_i c_i\, |o_i\rangle \otimes |\lambda o_i\rangle_{\text{app}}, \tag{97}$$

where the apparatus pointer is displaced by λo_i when the system is in eigenstate $|o_i\rangle$. Substituting into the substrate-projection picture: the substrate configuration C now feels a modified tension field through its coupling to the apparatus, with outcome-dependent threshold

$$\tau_i(C) \;=\; \tau_0(C) + \kappa \left(\lambda o_i\right)^2, \tag{98}$$

where κ is the substrate-apparatus coupling coefficient (a P3 parameter). The configuration C crosses τ_i if and only if $|\Pi^{-1}(P_i\psi)|^2 > \tau_i(C)$, that is, the substrate-projected weight on outcome i exceeds the outcome-specific threshold. The outcome-dependent threshold τ_i of the abstract argument is therefore not a postulate but the structural consequence of the substrate's response to the apparatus coupling (96).

10.3.2.2. Closure (2): Decoherence inheritance. In macroscopic measuring devices, the apparatus is entangled with environmental degrees of freedom (the bath), and the apparatus-pointer states $|\lambda o_i\rangle_{\text{app}}$ rapidly become quasi-orthogonal on the decoherence timescale $\tau_D \sim \hbar/(k_B T)$ for thermal bath temperature T. Environment-induced superselection (Zurek [106]) selects the pointer-eigenbasis $\{|o_i\rangle\}_i$ as the einselected basis on which off-diagonal density-matrix elements vanish. In PST: the substrate threshold field τ_i inherits this einselected basis without modification. The substrate configuration C couples to the apparatus through the einselected pointer states, so the threshold-crossing event of P3 is sharply defined on the same basis as standard quantum measurement. PST makes no claim of independent decoherence theory; the standard einselection mechanism applies verbatim. This is the meaning of "decoherence inheritance:" PST extends standard quantum measurement by adding the substrate-projection layer, without rewriting the decoherence machinery.

It is worth being exact about what is being inherited, because the division of labour between the three closures depends on it and is easy to misread in PST's disfavour. Decoherence does two things and conspicuously does not do a third. It selects the pointer basis, and it destroys interference in that basis. It does *not* make an outcome definite. The evolution is unitary, so the global state stays pure, and what is left on the system is an *improper* mixture: a reduced density matrix numerically identical to a proper statistical mixture, while the formalism carries no fact about which outcome obtained [206]. Computation 146 verifies both halves directly: global purity 1 throughout, and agreement with a hand-built proper mixture to machine precision under every local observable tried. This is standard quantum mechanics, not a PST claim; it is the reason the measurement problem survives decoherence at all.

The three ingredients therefore do disjoint work. The configuration-basis rule on the Bernoulli measure, its form fixed by Gleason's theorem, supplies the *probabilities* (§10.2); einselection supplies the *basis* and the loss of interference; and the threshold crossing of P3 supplies the *definiteness*, and with it a

single world. The crossing is not a restatement of decoherence in substrate vocabulary. It performs precisely the step decoherence cannot, and it does so without a collapse postulate and without branching.

That also settles a question the projected phenomenology raises later. Computation 144 finds that no fringe-visibility measurement separates a world in which the crossing selects from one in which it does not, and it would be easy to read that as the crossing doing no work. The reading is backwards. Improper and proper mixtures are locally indistinguishable *by construction* (Computation 146), so no statistic gathered on the system could have separated them. The absence of a signature is what the formalism requires, not evidence of an idle layer.

10.3.2.3. Closure (3): threshold-crossing uniqueness. We state and prove the substrate-level theorem that anchors the collapse rule.

10.3.2.4. Theorem (threshold-crossing uniqueness). *Let $\hat{O}$ be a self-adjoint observable with spectral decomposition $\hat{O} = \sum_i o_i P_i$ on the projected Hilbert space $\mathcal{H} = L^2(M, d^4x)$. Under the measurement coupling* (96) *and the outcome-dependent threshold* (98)*, the substrate threshold-crossing event of P3:*

(U) *Selects exactly one outcome o_i: for the substrate configuration C, the set $\{i : |\Pi^{-1}(P_i\psi)|^2 > \tau_i(C)\}$ is a singleton (or empty if no crossing occurred).*

(I) *Is irreversible: the post-crossing configuration $C' \in \mathcal{S}_i$ does not revert to $C \in \mathcal{P}(D) \setminus \mathcal{S}_i$ under the substrate's directed tension dynamics of P2–P3.*

Proof of (U). By construction, the eigenspaces $\{\mathcal{H}_i\}_i$ are orthogonal: $\mathcal{H}_i \perp \mathcal{H}_j$ for $i \neq j$, so $|\Pi^{-1}(P_i\psi)|^2 + |\Pi^{-1}(P_j\psi)|^2 \leq \|\psi\|^2$ for $i \neq j$. Together with the strong-coupling condition $\kappa\lambda^2 \gg \tau_0$, the outcome-dependent thresholds $\tau_i(C) = \tau_0(C) + \kappa(\lambda o_i)^2$ are *outcome-separated*: for $i \neq j$ with $|o_i| \neq |o_j|$, the thresholds τ_i, τ_j differ by $\kappa\lambda^2(o_i^2 - o_j^2)$. In the strong-coupling limit, this

gap exceeds any single $|\Pi^{-1}(P_i\psi)|^2$ that has not yet been concentrated on a specific outcome by the einselected basis (Closure 2 above). Therefore at most one τ_i is crossed by $|\Pi^{-1}(P_i\psi)|^2$, delivering uniqueness. In the degenerate case $|o_i| = |o_j|$ with $i \neq j$, the einselection of Closure (2) selects a pointer-basis projection that re-separates the thresholds before the crossing, so the same argument applies on the einselected subspace. □

Proof of (I). The substrate's asymmetric tension structure of P2 (the directed modal threshold of P3) is encoded by the parity-violation mechanism of Computations 8 and 11 (the $T(C)$- configuration breaks the substrate-side parity transformation of $\mathrm{SU}(2)_L$ irreversibly). Concretely, the threshold-crossing move $C \to C' \in \mathcal{S}_i$ involves a directed flow on the substrate's tension landscape: the configuration C is at a threshold-saddle that flows to C' under the P2 asymmetric tension. The reverse flow $C' \to C$ would require crossing the same threshold-saddle in the opposite direction, which requires a sign reversal of the P2 directed tension. The P2 asymmetric primitive, "modally distinguishable in a directed way", forbids the sign reversal at the substrate level. Therefore $C' \to C$ is forbidden in the substrate's directed dynamics, and the threshold-crossing is irreversible.

The directed-tension argument generalises the parity-violation mechanism of Computations 8 and 11 to arbitrary observables: the $\mathrm{SU}(2)_L$ parity violation is the prototype example ("handedness" of the weak charge), but the same asymmetric P2 structure forbids the reversal of any threshold-crossing event. The measurement postulate's irreversibility, standardly justified by appeal to thermodynamic entropy increase, is delivered in PST by the substrate-level asymmetric tension. □

Closures (1)-(3) above complete the measurement postulate derivation. The state-update rule (collapse) is a structural theorem from: (i) the von Neumann measurement coupling (96) which sets up outcome-dependent thresholds (98), (ii) the decoherence-einselection inheritance which sharpens the pointer basis,

and (iii) the uniqueness + irreversibility theorem (U), (I) above which selects a single outcome. Together with the Born probability of equation (95), this is the full measurement postulate, both probability rule and state-update rule, structurally delivered by PST's substrate machinery with no additional axiom beyond the assumed unitary propagation between crossings.

10.3.2.5. Closure (4): agreement with the Born rule. Closures (1)–(3) settle the state-update rule. They leave a consistency question that should be stated rather than assumed, because the probability rule and the state-update rule are derived by different routes and have to agree. The Born rule comes from the configuration-basis rule on the Bernoulli measure (§10.2), Gleason's theorem fixing its form, and depends only on the projected weights. The crossing rule, by contrast, selects an outcome by comparing those weights against a threshold. Nothing so far forces the two to give the same answer.

Write the question precisely. For a fixed state ψ the weights $w_i = |\Pi^{-1}(P_i\psi)|^2$ are fixed numbers, so the run-to-run randomness is not in them. It is in the substrate configuration C, entering through the tension magnitude $|T(C)|$ that §4.1 tests against the threshold. Selection frequencies are therefore

$$\mathbb{P}(\text{outcome } i) \;=\; \mathbb{P}_{C\sim\mu}\big(\, C \text{ crosses } \tau_i \,\big), \tag{99}$$

and agreement with Born requires this to equal w_i. Computation 143 shows that it does *under a condition it identifies and does not discharge*: the tension variable must be *uniform in its own quantile* across configurations. Given that, $\mathbb{P}(\tau_0$ below the outcome-i margin$) = w_i$ exactly.

That condition is not available for free, and an argument that it was is withdrawn. The argument ran: every derivation downstream of (17) reads T only through the *ordering* of the magnitudes, Computation 140 verifies invariance under monotone reparametrisation, the probability-integral transform is such a reparametrisation, so uniformity is a choice of scale rather than an assumption. The step that fails is the first. (98) does not compare T against T; it compares $\tau_0(C)$ against the Born weights w_i, which come from the quantum

state and carry their own normalisation $\sum_i w_i = 1$. That is a *cardinal* read: it needs to know where w_i sits on the tension scale, not merely how tension values order among themselves. Applying a monotone ϕ changes $\mathbb{P}(\text{outcome } i)$ from $F(w_i)$ to $F(\phi^{-1}(w_i))$ for every $\phi \neq \text{id}$, and the results are not even probability vectors: with $w = (0.5, 0.3, 0.2)$ and τ_0 uniform, $\phi(x) = x^2$ gives $(0.71, 0.55, 0.45)$ summing to 1.70, and $\phi(x) = x/2$, which is *affine*, gives $(1.00, 0.60, 0.40)$ summing to 2.00.

Computation 140's invariance looks applicable because it co-transforms *both* sides, setting $\tau' = \phi(\tau)$ and checking $\text{sign}(\phi(|T|) - \phi(\tau))$; that hypothesis is silently dropped here, where only one side transforms. Its own scope note says it audits the reads identified in the body, and this one was never among them. The dilemma is clean: if the crossing rule is a genuine read of T, then T is not free up to monotone reparametrisation and uniformity is a substantive constraint on the substrate; if T really is free, the rule fixes no frequencies at all. Either way *Closure (4) does not close on its own*, and the residue is a condition, which the rest of this subsection reduces to the non-degeneracy of T's contingent content and P3's typicality supplies (Appendix C).

What *is* established, and is worth having, is that the condition is *satisfiable*. The obstruction is ties. Configurations sharing a tension value cannot be separated by any order-preserving relabelling, so the largest such tie puts a floor on how uniform the variable can be made. Computation 145 measures it on the scalar exemplar T_{ex} of (17), by exact enumeration of all $2^{|D|}$ configurations. For generic distinction structures the floor is already small at $|D| = 14$ and falls as $|D|^{-1.5}$, so it is negligible for any substrate large enough to carry an apparatus. For *regular* (D, δ) the floor is larger at those sizes, T_{ex} there depending on C only through $|C|$.

A tempting shortcut fails (Computation 180). It would say: that is the same sub-class (17) already excludes, since T_{ex} fails complement-asymmetry on regular graphs, so the two exclusions coincide. *Both halves of that are wrong.* (16) is a condition on v, not on δ: its genuine content is the non-balance $\left|\sum_a v(a)\right| > 0$, as stated at §3.4, so it excludes no distinction graph whatever.

And the regular graphs are not the failing class in any case, their floor vanishing as $|D|^{-0.5}$ *on the exemplar* T_{ex}, slower than generic but vanishing; *that exponent is an artefact of integrality and is corrected* (Computation 187), the complete graph's atom falling exponentially under a generic admissible T rather than as a square root, so those families pass for a better reason than was recorded. *What does fail* is a different property: a *bounded active part*. A star has one centre and K_k with isolated vertices has a fixed clique, so in both T_{ex} takes O(1) distinct values however large $|D|$ grows; the star keeps an atom of exactly $1/2$ at every $|D|$, half of all configurations omitting the centre and tying at zero, which floors the distance from uniform at $1/4$ under every monotone reparametrisation. Both satisfy (14) and (15), checked exhaustively, and (16) cannot touch them, so both are *admissible*. *The consequence for the register*: Born agreement is jointly a condition on the law of T *and* on δ, and the theory constrains δ not at all in the relevant respect, so discharging it looked to need a genericity constraint on the distinction graph rather than a calculation. Nothing downstream is touched, and the generic case Computation 145 establishes is confirmed rather than disputed.

The genericity constraint is not needed: the graph half is a theorem (Computation 190). What Computation 183 measured, that no reduced family keeps the forced atom finite, is proved in two lines. Any matching of m pairwise disjoint distinctions gives m independent events of probability $\frac{1}{4}$ under μ, each the event that both ends of one distinction are present, and a distinction-free configuration avoids all of them, so the forced atom is at most $(3/4)^m$. Reducedness, P1(iii)'s consequence that no two properties share a distinction set (Computation 184), forces the matching large: the properties outside a maximal matching form an independent set, each has its distinction set inside the $2m$ matched properties, and injectivity bounds their number by 4^m, so $|D| - 2m \leq 4^m$. Together, for every reduced (D, δ) and every admissible T, the atom at zero is below $\max\big((3/4)^{|D|/4}, (|D|/2)^{-0.21}\big)$, a universal bound that vanishes in $|D|$ and is 10^{-17} at substrate scale. The bound is checked exactly against every family the earlier computations tabulated, and it is

universal rather than sharp: the subset-graph family that Computation 183 built as the extremal reduced case has an atom that falls as $|D|^{-1}$ against the bound's $|D|^{-0.21}$. With Computation 187, every atom but the forced one belongs to T, so what remained was a condition on the law of T alone, that its contingent content does not tie. Computation 193 shows that condition to be one face of the non-degeneracy of the content, non-balance being the other, and Computation 195 that the maximum-entropy law on the content has a density, under which both faces hold almost surely. P3's typicality is read as reaching the content of T (§1.7), so the condition is supplied on the footing of A3 and has left the structural front; Appendix C carries it as a conditional theorem.

One scope limit, stated plainly. Equation (98) cannot deliver uniqueness and Born statistics at the *same* coupling. Raising $\kappa\lambda^2$ raises every threshold, which is exactly what drives the probability of two simultaneous crossings to zero and is what clause (U) rests on. But pushed far enough it also orders the outcomes by margin, and the survivor becomes the same one every run. Born statistics live at the weak-coupling end of that axis and uniqueness at the strong-coupling end (Computation 143). The measurement postulate as assembled here is therefore consistent in the regime where the einselected pointer basis of Closure (2) does the separating, with (98) supplying irreversibility rather than the selection statistics. Locating that regime sharply is open.

10.3.3. Why this is not the many-worlds picture

Many-worlds interpretations of quantum mechanics also identify the measurement postulate with a branching of the wavefunction at the measurement event. PST differs in one essential respect: PST's threshold-crossing of P3 is a *single* irreversible event delivering one configuration, by the directed-threshold mechanism that already structurally underwrites parity violation (Computations 8, 11). There is no branching into multiple co-existing post-measurement states; there is a single post-measurement state $C' \in \mathcal{S}_i$ for the actually-realised outcome i. The PST collapse rule is therefore intrinsically one-world.

10.3.4. Why this is not an objective-collapse theory

The rule above shares an aim with the objective-collapse programme. GRW and its continuous-spontaneous-localisation descendants [195, 196] also replace the measurement axiom with a dynamical law rather than an interpretation, and they are the closest existing relatives of the rule set out above. They differ from PST on the one point that is empirically live.

Those models add a *spontaneous* localisation rate to the Schrödinger evolution, scaling with particle number or mass, so that an isolated superposition localises on its own once it is large enough. PST has no such term, and its absence is structural rather than a choice of parameter. The threshold a configuration must cross is $\tau_i(C) = \tau_0(C) + \kappa(\lambda o_i)^2$ of (98), in which λ is the displacement of an apparatus pointer. For a system coupled to nothing, λ has no value and the outcome-dependent thresholds τ_i do not exist; what remains is the unmodified modal threshold τ of P3, which is not outcome-dependent, so there is nothing to cross and no member of $\{\mathcal{S}_i\}_i$ to select. That is the argument already run for the undetected slit in §10.4, here stated in general: in PST collapse is triggered by coupling, not by size.

Three consequences follow, in increasing order of exposure.

- *No size limit on coherent superposition.* Nothing in the collapse machinery of this chapter depends on mass or particle number, so a superposition isolated well enough survives at any scale. Matter-wave interferometry has carried molecules beyond 25 kDa [197], and each such result bounds the objective-collapse parameter space from one side while leaving PST where it was. This is a contrast rather than a test: the two accounts agree wherever the postulated collapse rate is too small to resolve.

- *The falsifier is a collapse without a coupling.* A localisation observed in a system coupled to no apparatus, at a rate rising with mass, refutes the rule of this chapter outright. No term in τ_i could produce it, and no parameter could be adjusted to accommodate it afterwards.

- *PST cannot retreat by adding a rate, and this is where the exposure is*

sharpest. The one rate-like quantity in the theory is the equilibrium density j_τ of threshold-crossed configurations, and §16.5 defines it as a saturation density against the configuration-space measure of P3, the substrate having no spatial extent against which a proper-volume rate could be defined. That definitional choice is exactly what forces $w = -1$ rather than $w = 0$: a per-emergent-volume creation rate gives Hoyle–Bondi–Gold steady state, not a cosmological constant. A spontaneous collapse rate is a per-volume, per-time quantity of precisely the excluded kind.

The third point is worth stating plainly, because it closes an escape route before the theory has need of one. Were spontaneous localisation observed, PST could not absorb it by fitting a rate the way a collapse model fits its localisation parameter; it would have to give up the pre-spatial framing of the substrate, and with it the derivation of $w = -1$ in §16.5. *The collapse rule and the dark-energy equation of state are exposed together*, and that is a sharper commitment than either looks on its own.

10.4. The canonical experiments

The measurement postulate assembled in §10.3 has so far been exercised only on abstract observables. This section puts it to work on the experiments a reader arriving from physics tests a foundational proposal against first: the double slit, Wheeler's delayed-choice arrangement together with the quantum eraser, and the Stern–Gerlach magnet. One caveat governs all three and is stated here once rather than three times. PST supplies no propagation rule of its own between threshold-crossing events. It derives the arena, the projected Hilbert space of §10.2 and the canonical commutation relations of §10.6, but it nowhere derives that the projected state evolves between crossings under a one-parameter unitary group. That evolution is assumed below exactly as it is assumed in §10.3. In consequence nothing in this section is a new prediction. The probabilities are the Born probabilities of §10.2, so every number PST returns for these experiments is the number standard quantum mechanics returns already, and no experiment of this family can

distinguish the two. What is gained is an account of the measurement event itself, single, irreversible and one-world, with no observer in a privileged role and no branching. PST's one laboratory-scale falsifiable departure from standard physics remains the d^{-6} Casimir correction of Chapter 15.

10.4.1. The double slit and which-path information

The double slit poses two questions at once: why an interference pattern accumulates from individually localised detections [188] when the two paths are not distinguished, and why that pattern is absent when they are. Standard quantum mechanics answers the first by superposition and the second by the measurement postulate, and leaves the boundary between the two regimes to be drawn by hand, typically by an appeal to what counts as an observation. PST draws it structurally, and the criterion is the one already in force in P3: the fringes survive exactly when no threshold-crossing event occurs at the slits, and are destroyed exactly when one does.

Consider first the apparatus with no which-path detector. The substrate configuration C passes the slit plane with nothing coupled to its tension structure there. No apparatus pointer momentum enters a coupling of the form (96), so the threshold in force is the unmodified $\tau_0(C)$ of (98), the modal threshold τ of §4.1, which is not outcome-dependent. With no outcome-dependent threshold there is no crossing; with no crossing there is no selection; and with no selection no member of the partition $\{\mathcal{S}_i\}_i$ has been singled out at the slits. The projected wavefunction arriving at the screen therefore retains the contributions of both slit regions. The first threshold-crossing of the experiment is the one at which the substrate couples to the screen, whose observable is position rather than path, and whose eigenspaces each receive both contributions. The Born rule of (94) evaluated on that state returns $|\psi|^2$ with both contributions present, which is the fringe pattern. Interference and pointlike detection are not in tension in this account. The first is a property of a state that has not crossed a threshold, the second is a property of the crossing event, which by P3 delivers one configuration and hence one localised

mark; the pattern is the accumulated statistics of many such single crossings.

Now install a which-path detector at the slits. Its coupling is of the von Neumann form (96), with $\hat{O}$ the path observable and its two projectors P_1, P_2; the substrate's tension structure feels the apparatus and the threshold becomes the outcome-dependent τ_i of (98). The uniqueness half of the threshold-crossing theorem of §10.3, clause (U), then applies at the slit plane: at most one τ_i is crossed, so exactly one path is selected. The einselected pointer basis of Closure (2) is what makes the two path outcomes separable in the degenerate case where the labelling assigns them equal magnitude. The irreversibility half, clause (I), underwritten by the directed parity-violating mechanism of Computations 8 and 11, fixes that selection: the post-crossing configuration $C' \in \mathcal{S}_i$ does not revert. The state reaching the screen therefore carries the contribution of the selected slit alone, and the accumulated distribution is the sum of two single-slit distributions. Nothing has erased a pattern that formed earlier. The pattern never formed, because a configuration was already selected at the slits.

What this buys is that the distinction between measurement and no measurement is objective and physical. The discriminator is whether a coupling modifies τ into an outcome-dependent τ_i, which is a question about the substrate's tension structure and the device it is coupled to, and it has a determinate answer independent of anyone's knowledge, attention or existence. A which-path detector whose record is never read destroys the pattern all the same, because the crossing occurred at the slits and the irreversibility of clause (I) does not wait upon a reader. An observer who watches the slits without coupling anything to τ removes no fringe. The anthropocentric residue that the measurement postulate carries in its standard presentation is therefore absent here, and so is the alternative of keeping every outcome. By the argument of §10.3 the crossing is single and the collapse intrinsically one-world, so the surviving contribution is the only one there is, and not one branch among co-existing others.

Partial which-path information sits between the two cases, and the limits of what the machinery settles there are worth stating plainly. The threshold separation on which the proof of clause (U) rests is $\kappa\lambda^2(o_i^2 - o_j^2)$, which falls away with the coupling strength λ, and that proof is given in the strong-coupling limit $\kappa\lambda^2 \gg \tau_0$ only. A weak which-path coupling accordingly leaves the thresholds too little separated for a crossing to be forced at the slits on every passage.

It is tempting to read a prediction into that gap, and for a time this book did: if only some passages cross, the pattern should degrade rather than vanish, and the rate of degradation might be a signature PST could be tested by. Computation 144 closes that reading, and against it. Were a crossing at the slits a physical selection, only the non-crossing fraction would interfere, and the Englert–Greenberger–Yasin relation [203, 204] would be violated *downward*, most strongly at weak coupling, which is the regime in question. Atom-interferometric tests of that relation see no such deficit [205]. What the gap in clause (U) leaves open is therefore filled by the decoherence of Closure (2), which PST inherits verbatim rather than rederiving: the crossing is redundant with it, and the quantitative complementarity between fringe visibility and path distinguishability is the standard one, not something PST modifies (§19.9).

That is a narrower claim than the one this subsection might have made, and it is worth being explicit about what it costs. In the partial-coupling regime PST adds no new phenomenology; it reproduces standard quantum mechanics, as it must, since recovering quantum mechanics is what the projection is for. The claim that survives is structural rather than quantitative: the familiar phenomenology of the double slit follows from P1–P3 with no separate measurement axiom and no observer, and the line between the two regimes is drawn by the substrate itself rather than by a reader.

10.4.2. Delayed choice and the quantum eraser

Wheeler's delayed-choice arrangement [189] sharpens the measurement question by moving the experimenter's decision late. A photon enters an interferometer, and only after it has entered is it settled whether the recombining element is in place, so that the two paths interfere, or absent, so that the path taken is recorded. The statistics follow the configuration finally chosen. The apparent difficulty is that the photon seems to have had to behave in one way or the other at the entrance, and that a choice made afterwards then fixes which. Stated in that form, the experiment appears to exhibit a later event acting upon an earlier one.

The premise is what fails, and the collapse rule of §10.3 says precisely where. A measurement in PST is the crossing of an outcome-dependent threshold τ_i, and τ_i does not exist until the substrate's tension structure couples to the measuring device. Before that coupling there is the modal threshold τ of P3 and nothing that refines it: the partition $\{\mathcal{S}_i\}_i$ of $\mathcal{P}(D)$ is defined by the eigenspaces of the observable $\hat{O}$, and $\hat{O}$ is whatever the apparatus, in the configuration it finally takes, measures. An interferometer with its recombining element in place and one without it do not measure the same observable; they induce different partitions of the configuration space, and (98) assigns them different thresholds. There is therefore no earlier assignment of C to some $\mathcal{S}_i$ for the later choice to overturn. Exactly one threshold-crossing occurs in a run of the experiment, and it occurs at the coupling which the completed apparatus defines. The theory has nothing to revise because it posited nothing that stood in need of revision.

This must be separated from a stronger claim which PST does not make and does not need. Nothing in the foregoing permits a later event to act upon an earlier one. Retrocausation would require a causal relation running from the choice back to something already settled, and neither term is available here. There is no earlier settled fact. And the causal relation is not fundamental in this theory but is coinstantiated with the metric (Chapter 5), generated by

the Lorentzian signature of the projection rather than presupposed by it. The substrate that carries C possesses no temporal order of its own, its primitives being distinction and tension, neither of which orders anything. So "the choice came later" is a statement about the emergent order in which the laboratory record is written, and within that order it is entirely correct. What it is not is a statement which dates the substrate-side transition and places it before the choice. The paradox is generated by taking it to do both at once, and it is that conflation, rather than the order of the laboratory events, which PST declines.

The quantum eraser [190] is usually presented as the sharper case, on the ground that which-path information once recorded is afterwards destroyed and the interference returns. In PST the first half of that description cannot be right. A threshold-crossing is irreversible, by clause (I) of §10.3 and the directed-tension mechanism which Computations 8 and 11 underwrite: the post-crossing configuration $C' \in \mathcal{S}_i$ does not return to $\mathcal{P}(D) \setminus \mathcal{S}_i$. Nothing an experimenter does afterwards un-crosses a threshold. What the erasing operation performs is a further measurement, upon a correlated partner system, whose own crossing is equally single and equally irreversible, and which thereby labels each already-recorded detection with a second outcome. Interference reappears when the record is sorted by that label [191]. It does not reappear in the record unsorted.

That distinction is not a technicality; it is the whole of the matter. The raw counts at the signal detector carry no fringes, and sorting cannot give them any, for the trivial reason that a partition of a record into subsets leaves the record itself untouched. Were it otherwise the unsorted counts would report the setting of a distant apparatus. The fringes live in the correlated subensembles, each picked out by a label which was written before any sorting was contemplated; selecting a subensemble alters no detection and creates none. The standing caveat of this section applies here with particular force, and is worth repeating in one line: nothing above is a prediction, and every number PST returns for a delayed-choice or eraser arrangement is the number

standard quantum mechanics returns already. What is gained is not empirical content but an account of the measurement event itself, set in a framework where the temporal order is derived rather than assumed.

10.4.3. Stern–Gerlach and the definite eigenvalue

The experiment first performed by Gerlach and Stern in 1922 [192] is the cheapest place to put the collapse rule to work, because the classical expectation and the observed result differ in the plainest possible way. A neutral atom carrying a magnetic moment of magnitude μ through an inhomogeneous field feels a transverse force $\partial_z(\mu_z B_z)$, so its deflection is proportional to μ_z. Classically the moment is a vector free to point anywhere on the sphere, μ_z takes every value in $[-\mu, +\mu]$, and a beam of randomly oriented atoms must paint a continuous band on the screen. It paints two spots. The discreteness is not itself a mystery for PST, because it is already present in the projected Hilbert space before any measurement is considered. The orthogonal decomposition into eigenspaces $\{\mathcal{H}_i\}_i$ over which the Born rule of §10.2 is stated has, for the moment of a spin carrier, two summands rather than a continuum. The Mosco limit delivers a four-dimensional Lorentzian $M = \mathbb{R} \times S^3$ whose local Clifford algebra is $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ (§13.7, Computation 66), and the substrate's odd-parity sector projects to fields that are sections of the associated spinor bundle (§11.2). The two-valuedness of μ_z is a corollary of that module rather than an added assumption, although, as the closing paragraph of this subsection records, the present text nowhere draws the corollary explicitly.

The question PST is actually needed for is the second one: why a given atom arrives at one of the two spots rather than in some superposition of both. Here the inhomogeneous field is the tension-structure coupling. In the abstract argument the apparatus replaces the global modal threshold τ with an outcome-dependent τ_i through equation (98); in a Stern–Gerlach magnet that substitution is realised by an ordinary field gradient. The observable in the von Neumann coupling (96) is μ_z and the apparatus pointer is the atom's transverse momentum, which the gradient displaces in proportion to

it. PST adds no force to this: the deflection is the electromagnetism of the projected Standard Model Lagrangian (195) and nothing else. What PST adds is that once the coupling is on, the modal threshold is no longer one global number. It is worth noticing which branch of the uniqueness theorem the experiment falls into. The two eigenvalues are $\pm\mu$, so $o_+^2 = o_-^2$ and the two thresholds of equation (98) coincide: uniqueness is therefore not delivered here by the strong-coupling threshold gap but by the degenerate clause of the proof of clause (U), in which the einselection of Closure (2) re-separates the thresholds on the pointer basis before the crossing. In the laboratory that clause is the screen. The magnet alone does not complete a measurement; the magnet correlates μ_z with a macroscopic position, and it is the irreversible registration of that position which fixes the basis the substrate threshold field then inherits.

Given the basis, the rest is the collapse rule applied verbatim. The configuration C crosses one of the two thresholds and not both, by clause (U), and it does not cross back, by clause (I): the directed tension of P2 forbids the sign reversal that a reverse crossing would require. The atom therefore lands at one place on the screen, with probability $\|P_\pm\psi\|^2$ by equation (95), which for an unpolarised beam is one half on each spot. The emphasis belongs on the word one. The outcome that was not selected is not realised on some other screen or in some other branch; it is simply never instantiated, which is the content of the one-world reading argued in §10.3. A single irreversible directed threshold crossing has delivered a definite eigenvalue out of what the classical account expected to be a continuum, and it has done so without a second postulate beyond P3.

Three magnets in series, oriented z, then x, then z again, make the shape of the account sharper, because the third magnet recovers both outcomes and a naive reading of irreversibility would forbid it. The first crossing leaves the substrate in a configuration $C' \in \mathcal{S}_+^{(z)}$. The second magnet measures a different observable, whose eigenspaces induce a different partition $\{\mathcal{S}_i^{(x)}\}$ of $\mathcal{P}(D)$, and because μ_x and μ_z do not commute the two partitions are incompatible: no

cell of one sits inside a cell of the other. Clause (I) says that the z crossing cannot be undone. It does not say that the configuration carries a permanent z label. What the first crossing fixed was membership of $\mathcal{S}_+^{(z)}$ and nothing further, and the x magnet installs a fresh tension structure whose thresholds are separated on a basis bearing no relation to the first. The crossing at the second magnet is thus a new selection, weighted $|\langle x_\pm|z_+\rangle|^2 = 1/2$, and the third magnet finds both z outcomes at equal weight. On the substrate account the re-randomisation is expected rather than anomalous, and for a structural reason rather than a statistical one: a threshold crossing records an eigenspace membership for the observable actually measured, and PST provides nothing else in which a record of it could survive an incompatible measurement.

What this subsection does not do should be stated as plainly as what it does. The standing caveat of the section applies here too, and in a concrete form: the atom has been treated as propagating coherently through the field region and between magnets, and no such propagation rule is derived anywhere in this work. Beyond it there are three further debts, all of them specific to this experiment. First, the two-valuedness itself is available but undrawn: $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ and the existence of the spinor bundle are derived under P1–P3, yet no passage in this book takes the step from the spinor module to the statement that a spin projection has exactly two eigenvalues. Second, nothing in P1–P3 has been shown to force the silver atom's moment to be carried by a single unpaired electron; that is atomic shell structure, and it is imported. Third, the magnitude of the splitting is imported as well: the moment of a Dirac particle follows in principle from the minimally coupled Dirac term of equation (195), but the non-relativistic reduction that extracts it is performed nowhere in this work. The number of spots and the size of the splitting are therefore inherited rather than computed. What the subsection establishes is the narrower point: a definite eigenvalue, rather than a superposition, is what P1–P3 deliver at the screen, and it costs no measurement postulate and no branching.

10.5. Composite systems, the reduced state and no-signalling

Everything above treats a single system. The state space is $L^2(\mathcal{P}(D), \mu)$, the observable is one self-adjoint $\hat{O}$ on it, and the threshold-crossing of P3 selects one of its eigenspaces. That is enough for the Born rule, for the collapse rule and for the three experiments of §10.4, none of which needs a notion of "two systems". It is not enough for the correlation experiments, which are about two wings and about what one wing's setting does to the other's statistics. This section supplies what they need, and the supply is unusually cheap: the substrate's measure is already a product measure, so the tensor structure is a theorem rather than an axiom.

10.5.1. The bipartition is of the property domain, not of space

A composite system in standard quantum mechanics is two systems in two places, and the tensor product is written down to match. PST cannot start there, because at the substrate level there is no space to be in: $M = \mathbb{R} \times S^3$ is an output of instantiation (Chapter 5), and the only primitive available is the property domain D with its distinction relation. What the substrate does offer is a way of splitting D itself.

Let $D = D_A \sqcup D_B$ be a partition of the property domain into two disjoint non-empty parts. Nothing in P1–P3 privileges one such partition over another: δ is symmetric and irreflexive and treats all pairs alike, which is the same even-handedness that forced the Bernoulli measure in §2.5. A bipartition is therefore not a structure the substrate possesses but a division an experimenter imposes, by building two apparatuses that couple to disjoint sets of properties. "Subsystem" in PST means *a part of the property domain*, and that is the whole of what it means.

10.5.2. Theorem (tensor factorisation over a bipartition)
Let $D = D_A \sqcup D_B$ and let $\mu = \bigotimes_{a\in D} \mathrm{Bern}(1/2)$ be the canonical measure of §2.5. Write $\mu_A = \bigotimes_{a\in D_A} \mathrm{Bern}(1/2)$ and μ_B likewise. Then the map

$$\beta\colon \mathcal{P}(D) \longrightarrow \mathcal{P}(D_A) \times \mathcal{P}(D_B), \qquad \beta(C) = (C \cap D_A,\ C \cap D_B) \tag{100}$$

is a measure-space isomorphism carrying μ to $\mu_A \otimes \mu_B$, and it induces a unitary

$$L^2(\mathcal{P}(D), \mu) \ \cong\ L^2(\mathcal{P}(D_A), \mu_A) \otimes L^2(\mathcal{P}(D_B), \mu_B) \ =\ \mathcal{H}_A \otimes \mathcal{H}_B. \tag{101}$$

Proof. That β is a bijection is immediate: a subset of D is determined by its trace on D_A and its trace on D_B, and any pair of traces arises from exactly one subset, namely their union. Under the identification $\mathcal{P}(D) \cong \{0,1\}^D$ it is the canonical splitting of a product index set, $\{0,1\}^D = \{0,1\}^{D_A} \times \{0,1\}^{D_B}$.

For the measure, μ is by construction a product over the index set D, so it factors over any partition of that index set. On the cylinder sets of (9), which generate the σ-algebra, take finite $C_0 \subseteq D$ and split it as $C_0 = C_0^A \sqcup C_0^B$ with $C_0^A = C_0 \cap D_A$:

$$\begin{aligned} \mu\big(\{C : C_0 \subseteq C\}\big) &= 2^{-|C_0|} \ =\ 2^{-|C_0^A|} \cdot 2^{-|C_0^B|} \\ &= \mu_A\big(\{C_A : C_0^A \subseteq C_A\}\big)\, \mu_B\big(\{C_B : C_0^B \subseteq C_B\}\big). \end{aligned} \tag{102}$$

The two measures agree on a generating π-system of finite total mass, so they agree (Dynkin), giving $\beta_*\mu = \mu_A \otimes \mu_B$.

The unitary is then the standard isomorphism $L^2(X \times Y, \mu_A \otimes \mu_B) \cong L^2(X, \mu_A) \otimes L^2(Y, \mu_B)$ for σ-finite measure spaces, which both are, being probability spaces. Explicitly, $\psi \mapsto \psi \circ \beta^{-1}$ carries the inner product (93) to the product inner product, and the elementary tensors $\phi_A \otimes \phi_B\colon (C_A, C_B) \mapsto \phi_A(C_A)\phi_B(C_B)$ span a dense subspace. Computation 142 verifies the factorisation numerically at finite $|D|$, to machine

precision. □

Two things about this deserve emphasis, because they are what make the rest work. The first is that **nothing was added**: the factorisation is a property of the measure the substrate already has, and that measure was forced by relabelling and complementation invariance, not chosen for convenience. Had μ been any non-product measure over D, (101) would fail and PST would have no composite systems at all. The second is that the factorisation holds for *every* bipartition simultaneously, which is the substrate-level statement of the fact that what counts as a subsystem is a choice of apparatus rather than a fact about the world.

10.5.3. The reduced state and local observables

With (101) in hand the standard apparatus is available and means what it usually means. An observable is *local to* A when its measurement coupling (96) involves only properties in D_A, so that it acts as $\hat{O}_A \otimes \mathbf{1}_B$. This is a definition, and it is the honest one: locality in PST is locality *in the property domain*, and it is prior to any spatial notion, which does not yet exist at the level where the definition is made.

For a state ρ on $\mathcal{H}_A \otimes \mathcal{H}_B$ the reduced state on A is the partial trace $\rho_A = \mathrm{Tr}_B\, \rho$, characterised by

$$\mathrm{Tr}\big[\rho\,(\hat{O}_A \otimes \mathbf{1}_B)\big] \;=\; \mathrm{Tr}\big[\rho_A\,\hat{O}_A\big] \qquad \text{for every } \hat{O}_A \text{ on } \mathcal{H}_A. \tag{103}$$

By the Born rule of §10.2, which is derived rather than assumed, the statistics of every observable local to A are therefore fixed by ρ_A alone. A state is *correlated across the bipartition* when ρ is not a product $\rho_A \otimes \rho_B$; this is a statement about the bipartition and says nothing about distance, there being none.

10.5.4. No-signalling across the bipartition

Now the question Bell's theorem makes urgent. Let the B wing measure $\hat{O}_B$ with spectral projectors $\{Q_j\}_j$. By §10.3 the coupling replaces the global threshold τ with the outcome-dependent τ_j on the B factor, and the threshold-crossing selects exactly one j, irreversibly. Over an ensemble of runs the outcome j occurs with Born weight $\mathrm{Tr}[\rho\,(\mathbf{1}_A \otimes Q_j)]$ and leaves the post-crossing state $(\mathbf{1}_A \otimes Q_j)\rho(\mathbf{1}_A \otimes Q_j)$ up to normalisation. The ensemble average over outcomes is therefore

$$\rho' \;=\; \sum_j (\mathbf{1}_A \otimes Q_j)\,\rho\,(\mathbf{1}_A \otimes Q_j), \qquad \mathrm{Tr}_B\,\rho' \;=\; \mathrm{Tr}_B\,\rho \;=\; \rho_A, \tag{104}$$

the second equality by cyclicity of the trace on the B factor together with $\sum_j Q_j = \mathbf{1}_B$. The reduced state on A is unchanged. Since by (103) every observable local to A has its statistics fixed by ρ_A, no choice of $\hat{O}_B$ at the B wing, and no outcome obtained there, alters any statistic available at A. **This is no-signalling, derived rather than asserted**, and its two ingredients are the tensor factorisation proved above and the Born rule of §10.2. Computation 142 checks it numerically on correlated substrate states, including states that violate the Clauser–Horne–Shimony–Holt (CHSH) inequality, and finds the A marginal invariant to machine precision under every choice of B setting.

The one-world reading of §10.3 is not weakened by this argument, but it does constrain what the argument says. In a single run the crossing at B is single and irreversible and there is a fact about which j occurred; (104) is a statement about the ensemble, as it is in standard quantum mechanics. What PST adds is that the ensemble is an ensemble of runs rather than a collection of co-existing branches, so the averaging in (104) is ordinary statistics over repetitions and not a sum over worlds.

10.5.5. Bell correlations and emergent spatiality

Bell's theorem [194] rules out local hidden-variable completions of quantum mechanics, where locality is a condition on how variables factorise across *spatial* separation. PST is not obliged to satisfy it in the form usually intended, and the reason can now be stated structurally rather than asserted. The substrate's own notion of a subsystem is a part of the property domain, by the bipartition above; spatial separation is not a substrate notion at all, since $M = \mathbb{R} \times S^3$ is an output of instantiation (Chapter 5). Substrate-level structure therefore carries no antecedent obligation to factorise along distinctions that only exist downstream of it, and a correlated ρ across a bipartition is not a puzzle but the generic case: product states are the special ones.

What must still hold, and now does hold as a theorem rather than as an expectation, is no-signalling: by (104) the marginal statistics at one part of the bipartition are independent of the setting at the other. PST's position is thus the familiar one for any theory taking spacetime to be emergent, but with the load-bearing step supplied: the correlations are not mysterious because they were never spatial in origin, and no operational influence is transmitted because the reduced state is provably untouched.

One link in this chain deserves the pressure a critic would put on it, and it is narrower than it first appears. Everything above is exact at the level of the property domain: the factorisation, the reduced state and no-signalling all hold across a bipartition $D = D_A \sqcup D_B$. The experiments, however, are performed at two *places*, so the chain is only about experiments if a spatial separation delivers such a bipartition.

It does, and the reason is the construction of §4.10 rather than anything new. There $\rho : D \twoheadrightarrow M$ is a quotient map and a spacetime point *is* a fibre of it, an equivalence class of properties rather than a container they sit in. A region $U \subseteq M$ is therefore already a set of properties, namely $\rho^{-1}(U) = \bigcup_{p \in U} p$, and fibres of a function are disjoint, so disjoint regions give disjoint subsets of D. Setting $D_A := \rho^{-1}(U_A)$ and $D_B := D \setminus D_A$ yields exactly the bipartition the

theorem above needs, and it does so without any further assumption: ρ is fixed by (D, δ) and the tension structure, not by the configuration, so D_A is a subset of D and not a configuration-dependent selection. The correspondence from places to parts of the property domain is thus immediate, and a reader should not be told the theory lacks it.

What is genuinely not established is narrower, and it is worth stating exactly, because it is what closing this line would have to address. Two things remain.

First, ρ is not unique. The freedom to partition D into fibres in more than one way consistent with the tension structure *is* PST's diffeomorphism invariance (§4.10), so D_A is defined relative to a choice of ρ. No ρ-independent statement can depend on that choice, and the metric relations that make the two wings spacelike-separated are preserved across admissible partitions by construction. But PST has not shown that "the A apparatus" is a ρ-independent designation, and until it does, the invariance is an expectation rather than a theorem.

Second, PST has no account of an apparatus. Identifying the A wing with $\rho^{-1}(U_A)$ says that a laboratory device occupying a region is realised by exactly the properties whose fibres lie in that region. That is close to definitional given that a point is nothing over and above its property-bundle. Of the rest, less is missing than it appears: whether such a set is spatially *compact* is not an extra fact about the geometry, because (30) already fixes the distance as $d(\rho(a), \rho(b)) = |T(\{a, b\}) - \tau_{\text{null}}|^{1/4}$, strictly monotone in the tension excess, so compactness in M and tightness in T are one condition stated twice (Computation 156). What PST genuinely does not derive is why T has such clusters, which is a fact about the configuration rather than about the structure.

The honest summary is therefore this. PST derives composite systems, reduced states and no-signalling outright, and the passage from a spatial bipartition to a bipartition of D is a corollary of the quotient construction rather than an assumption. PST never needs the naming of the A apparatus to be ρ-independent, the factorisation theorem being universally quantified over

bipartitions (Computation 151), and the compactness half reduces to (30). What it inherits is narrower than either: why the tension structure clusters at all. That is a question about what an apparatus is rather than about whether the geometry cooperates (§19.9.2).

10.6. Canonical commutation relations from the Boolean gradient

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

$$(\hat{x}_\mu\,\psi)(x) \;=\; x_\mu\,\psi(x), \tag{105}$$

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

$$\hat{p}_\mu \;=\; -i\hbar\,\Pi\big(\partial^{\mathrm{B}}_\mu\big), \tag{106}$$

where $\partial^{\mathrm{B}}_\mu$ is the Boolean (Efron–Stein) gradient (equation (20)) 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:

$$(\partial^{\mathrm{B}}_a\psi)(C) \;=\; \psi(C \cup \{a\}) - \psi(C \setminus \{a\}) \;\xrightarrow{|dx_a|\to 0}\; 2\,\partial_{x_\mu} f(x)\cdot \hat{e}_\mu, \tag{107}$$

where $\hat{e}_\mu$ 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 $\Pi(\partial^{\mathrm{B}}_\mu) = \partial_\mu$, 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 $\psi \in \mathcal{H}$:

$$\begin{aligned}[\hat{x}_\mu, \hat{p}_\nu]\,\psi &= -i\hbar\big(x_\mu\,\partial_\nu\psi \;-\; \partial_\nu(x_\mu\psi)\big) \\ &= -i\hbar\big(x_\mu\,\partial_\nu\psi \;-\; \delta_{\mu\nu}\psi \;-\; x_\mu\,\partial_\nu\psi\big) \\ &= i\hbar\,\delta_{\mu\nu}\,\psi. \end{aligned} \tag{108}$$

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 (107): 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 §7.8 (equation (58)) makes the limit (107) rigorous under A1–A7 (equicoercivity A7 is a theorem given A1, A2, A4 and A3, its equidistribution condition A3 supplied by P3's maximum-entropy sublimation, §7.1). 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 $L^2(M, \mathrm{d}^4x)$, 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 $C_c^\infty(M) \subset L^2(M, \mathrm{d}^4x)$ by standard elliptic theory [99]; 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 $\hat{p}_\mu$ is self-adjoint and densely defined. The canonical commutation relations are thereby established.

10.7. Quantum fluctuations

Near the modal threshold, where $\varepsilon(C) = |T(C)| - \tau \to 0$, the order parameter ψ is in the critical region: the two minima of the modal potential are about to

separate but have not yet done so. The potential near $\psi = 0$ is nearly flat in the direction of bifurcation, meaning that the order parameter requires almost no energy to shift significantly. Small fluctuations in the tension δT produce large fluctuations in the instantiated metric:

$$\delta g_{\mu\nu} \;\sim\; \eta_{\mu\nu} \cdot \frac{\delta T}{2\,|T - \tau_{\text{null}}|^{1/2}}\bigg|_{T \to \tau_{\text{null}}} \tag{109}$$

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 (30) 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 the settling of the order parameter from the near-threshold critical region into a stable minimum; when an apparatus is involved, that settling is the collapse of §10.3, the apparatus coupling shifting the threshold to an outcome-dependent τ_i and the configuration crossing it. Measurement in PST is therefore a physical interaction, but one that acts on the threshold rather than disturbing a pre-existing definite state: it is the process by which a near-threshold configuration is brought sufficiently far from τ that its projection becomes localised.

10.8. Summary

This chapter recovers quantum mechanics rather than predicting beyond it. The Born rule follows from the Bernoulli measure and Gleason's theorem on $L^2(\mathcal{P}(D), \mu)$ (§10.2); the full measurement postulate, state update included, follows from threshold-crossing, with the agreement condition on the law of T stated where it enters (§10.3). The canonical experiments, the double slit, delayed choice and the quantum eraser, and Bell's, come out as the postulate predicts (§10.4); composite systems carry the reduced state and no-signalling (§10.5); the canonical commutation relations are the image of the Boolean gradient under projection (§10.6); and quantum fluctuations are the critical region of the order parameter near threshold (§10.7). Measurement is a physical interaction on the threshold, not a disturbance of a pre-existing definite state.

11. Spin and Statistics from Boolean Exclusion

Matter particles obey a rule that keeps them out of each other's way: no two can be in exactly the same state. That rule is why solid things are solid. This chapter gets that rule out of the substrate rather than assuming it.

11.1. The fermionic sector: structure and the open problem

The excitations of the scalar sector (the order parameter, identified with the Higgs doublet $H \in \mathbb{C}^2$ in Chapter 13) are integer-spin fields: the Higgs scalar, the gauge bosons, and the graviton. Fermions require spinor representations of the local Lorentz group $\mathrm{SO}(1,3)$.

The spinor bundle exists. Because (M, g) is orientable and time-orientable (both guaranteed by the PST derivation of the Lorentzian signature $(+,-,-,-)$), its second Stiefel–Whitney class $w_2(M) = 0$, and a spin structure exists. The vierbein $e^a{}_\mu$ satisfying $g_{\mu\nu} = e^a{}_\mu e^b{}_\nu \eta_{ab}$ is a derived quantity (the square root of the projected metric), not an independent primitive. The 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 answered in the next paragraph and §11.2.

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 doublet's threshold vacuum manifold $\mathcal{V}_{\mathrm{EW}} \cong S^3 \cong \mathrm{SU}(2)$ (§13.4) 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 §11.2 below: the Boolean $\{0,1\}$ distinction structure is isomorphic to a fermionic Fock space via the canonical anticommutation relations (CAR) algebra, so fermionic excitations are not an additional postulate but are already present in the substrate.

11.2. Fermionic excitations: the CAR algebra identification

The question left open in the preceding section 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

$$c_a^\dagger c_a \;=\; n_a, \qquad \{c_a,\, c_b^\dagger\} \;=\; \delta_{ab}, \qquad \{c_a,\, c_b\} \;=\; 0, \tag{110}$$

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

$$c_a \;=\; \Big(\prod_{b<a}(1-2n_b)\Big)\,\sigma_a^-, \qquad \sigma_a^- \;=\; |0\rangle_a\langle 1|_a, \tag{111}$$

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

$$\mathcal{H}_D^+ \;=\; \bigoplus_{k\,\text{even}} \bigwedge^k \mathbb{C}^{|D|}, \qquad \mathcal{H}_D^- \;=\; \bigoplus_{k\,\text{odd}} \bigwedge^k \mathbb{C}^{|D|}. \tag{112}$$

Every state of $\mathcal{H}_D$ is antisymmetric under exchange of its *constituents*, even

and odd k alike, $\mathcal{H}_D$ being an exterior algebra. What distinguishes the two sectors is the fermion-parity grading: exchanging two *composites* of occupation k_1 and k_2 produces $(-1)^{k_1 k_2}$, which is $+1$ when both are even, and even-parity observables with disjoint support commute. So states in $\mathcal{H}_D^+$ project to bosonic fields under Π and states in $\mathcal{H}_D^-$ to fermionic fields, by composite exchange sign rather than by constituent symmetry.

The Bernoulli measure weights both sectors equally. Under $\mu = \bigotimes_a \mathrm{Bern}(1/2)$, 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 (110) require $c_a c_b = -c_b c_a$ for $a \neq b$, and the Jordan–Wigner string (111) 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 statistics half of the spin-statistics connection, that the projected fields anticommute, is therefore a structural consequence of the Boolean distinction primitive, not a separate postulate; the spin half is attached by the choice of spinor bundle rather than derived (§11.3). This is the *identification*; §11.3 promotes it to a formal theorem (CAR forced by Boolean exclusion, with the converse boson case and the Wightman-axiom development) and need not be re-read as the same point.

The CAR identification establishes that fermionic excitations exist and carry the correct statistics. The remaining fermion-sector results each need the gauge structure of Chapter 13 first, and are therefore taken up in Chapter 14 once it is available: the colour multiplicity $N_c = 3$ (§13.5, equation (163)); the generation count $N_{\mathrm{gen}} = 3$ (§14.2); the representation content of $\mathcal{H}_D^-$ under Π (§14.1); and the Yukawa hierarchy and Cabibbo–Kobayashi–Maskawa (CKM) mixing (§14.3). **Nothing in the present chapter depends on any of them**: spin and statistics come out of P1's Boolean exclusion alone.

11.3. Spin-statistics from Boolean exclusion (structural theorem)

The CAR identification of §11.2 already establishes that substrate excitations carry fermion statistics structurally. This section discharges the remaining formal work: it states that statement as a rigorous theorem (CAR algebra forced by Boolean exclusion, with the integer-spin boson case obtained by composition), contrasts it with the standard quantum field theory (QFT) spin-statistics theorem, and completes the functional-analytic development (substrate-to-continuum CAR convergence, the Wightman axioms, and the Dirac equation as substrate-derived). Recall the standard QFT theorem: half-integer-spin fields anticommute and integer-spin fields commute, derived on a relativistic field theory from Lorentz invariance, positive energy spectrum, locality (microcausality), and cluster decomposition. PST delivers the *statistics* half of that dichotomy from the Boolean distinction primitive of P1, without invoking the QFT axioms separately. *It does not deliver the spin half*: what is derived is that exclusion forces the CAR and that composition parity fixes statistics, and the spin label is attached by the choice of spinor bundle rather than derived, the book taking no step from the spinor module to a two-valued spin projection (§10.4.3).

11.3.1. Structural theorem (Boolean spin-statistics)

Let D be a finite distinction set and $\mathcal{P}(D)$ the substrate configuration space with Bernoulli measure $\mu = \bigotimes_a \mathrm{Bern}(\frac{1}{2})$. The Boolean structure on D, each site $a \in D$ being either occupied (1) or unoccupied (0) in a configuration $C \in \mathcal{P}(D)$, carries a canonical anticommutation-relations (CAR) algebra structure under the creation/annihilation operators

$$(c_a^\dagger \psi)(C) \;=\; \mathbf{1}[a \notin C] \cdot (-1)^{\#\{b<a : b \in C\}}\, \psi(C \cup \{a\}), \tag{113}$$

$$(c_a \psi)(C) \;=\; \mathbf{1}[a \in C] \cdot (-1)^{\#\{b<a : b \in C\}}\, \psi(C \setminus \{a\}), \tag{114}$$

satisfying $\{c_a, c_b\} = \{c_a^\dagger, c_b^\dagger\} = 0$ *and* $\{c_a, c_b^\dagger\} = \delta_{ab} I$ *(CAR),* not *the commutator relations* $[c_a, c_b^\dagger] = \delta_{ab} I$ *(CCR). Therefore the substrate-emergent excitations carry fermion statistics structurally from P1.*

Proof. The Jordan–Wigner construction (114) realises the Boolean $\{0,1\}$-exclusion principle as the constraint $(c_a^\dagger)^2 = 0$ (no double occupancy at a site): a single application of $c_a^\dagger$ to a configuration C with $a \in C$ already returns zero, and to C with $a \notin C$ the first application gives $C \cup \{a\}$ where a is then occupied, so the second application again returns zero. This is exactly the $c_a^\dagger c_a^\dagger = 0$ relation, equivalent under polarisation to $\{c_a^\dagger, c_a^\dagger\} = 0$. The cross-relation $\{c_a, c_b^\dagger\} = \delta_{ab} I$ follows from the Jordan–Wigner sign factor; the calculation is standard [105]. The corresponding commutator relations $[c_a, c_b^\dagger] = \delta_{ab} I$ would require $c_a^\dagger c_a^\dagger \neq 0$ (multiple bosons at a site allowed), which is incompatible with the Boolean exclusion of P1. □

11.3.2. Consequence: half-integer excitations are fermionic

Half-integer-spin elementary excitations on the substrate are fermionic by structure. Integer-spin excitations (photons from inner fluctuations of the gauge sector, composite bound states of even numbers of fermions) carry bosonic statistics by composition: an even-order composite has even fermion parity, so two of them exchange with sign $(-1)^{k_1 k_2} = +1$ and commute at disjoint support. *This is not the CCR.* For a pair operator $b = c_1 c_2$ one finds $[b, b^\dagger] = \mathbb{1} - n_1 - n_2$, so the canonical commutation relation is recovered only in the dilute limit where the constituent modes are empty; fermion bilinears close on a Lie algebra, not a Heisenberg algebra. The structural statement is therefore: *statistics of an excitation are fixed by its CAR-composition order; odd-order composites are fermionic, even-order composites are bosonic.* Note what that statement is indexed by: composition parity, not spin.

11.3.3. Relation to the standard QFT spin-statistics theorem

The QFT theorem assumes Lorentz invariance, positive energy, locality, and cluster decomposition, and *concludes* the spin-statistics rule. PST inverts

the logic: starting from the Boolean exclusion of P1, the CAR structure is forced, and those same relativistic properties follow downstream from the substrate projection Π and the Mosco-convergence emergence of $M = \mathbb{R} \times S^3$ (§7.1). The two arguments run from opposite ends and do not deliver the same statement: QFT derives a dichotomy indexed by *spin* from the relativistic arena, requiring its full apparatus; PST derives a dichotomy indexed by *composition parity* from the Boolean exclusion of P1 alone, with no further input from the spacetime emergence chain. The second is a real result and it is not the first.

11.3.4. Closure (1): CAR convergence

Extend the Mosco-convergence framework of §18.1 from the Boolean Dirichlet form to the substrate CAR algebra by identifying the substrate-side spinor degrees of freedom with the discrete creation/annihilation operators of equation (114). Concretely: choose a spinor bundle $\mathcal{S} \to M$ with structure group $\mathrm{Spin}(1,3)$ and trivialised over the Delaunay triangulation $\mathcal{T}_n$ of $\rho_n(D_n) \subset M$. Define the continuum spinor field operator

$$\psi(x) \;:=\; \lim_{n\to\infty} \sum_{a\in D_n} c_a \cdot \phi_a(x) \cdot s_a, \tag{115}$$

where $\phi_a(x)$ is the hat function on $\mathcal{T}_n$ centred at $\rho_n(a)$ (the same basis as in Closure (a) of §18.1.6) and $s_a \in \mathcal{S}_{\rho_n(a)}$ is a choice of spinor frame at site a, determined by the spin structure on S^3 (Computation 4 confirms uniqueness up to gauge). The limit (115) is taken in the strong-operator topology on the Gelfand–Naimark–Segal (GNS) Hilbert space of the substrate CAR.

11.3.5. Theorem (CAR Mosco convergence)

For $f, g \in L^2(M, d^4x) \otimes \mathcal{S}$ and any state $|\Phi\rangle$ in the GNS Hilbert space:

$$\sum_{a\in D_n} f(\rho_n(a))\, g(\rho_n(a)) \left\langle \Phi | c_a c_a^\dagger | \Phi \right\rangle \;\xrightarrow{n\to\infty}\; \int_M \langle f(x), g(x)\rangle_{\mathcal{S}} \left\langle \Phi | \psi(x)\psi^\dagger(x) | \Phi \right\rangle d^4x, \tag{116}$$

and the substrate CAR relations $\{c_a, c_b^\dagger\} = \delta_{ab} I$ *converge, on a Cauchy slice of the sublimation-derived timelike direction, to the equal-time continuum CAR* $\{\psi_\alpha(\mathbf{x}, t), \psi_\beta^\dagger(\mathbf{y}, t)\} = \delta_{\alpha\beta}\, \delta^{(3)}(\mathbf{x} - \mathbf{y})\, I$ *in the distributional sense, with the rate matching the Berry–Esseen rate of §18.1.6. The slice matters, and it is supplied* (Computations 186, 191, 192). Writing $\delta^{(4)}(x - y)$ instead, summing over sites realised into the four-dimensional M rather than into a slice, would be dimensionally wrong by one power of mass, since $\{\psi, \psi^\dagger\}$ carries [mass]^3 in four dimensions; it contradicts the free Dirac field, whose unequal-time anticommutator is $iS(x - y)$ rather than a delta; and it is incompatible with the spectral condition W2 stated below, since a constant Fourier transform cannot have support in the closed forward cone. Posing the relations slice-wise is what the substrate's lack of a preferred time makes non-trivial, and it is done in three steps. No site is realised *on* a measure-zero slice, so the relations are posed on a *slab* of thickness h_n whose site density per unit 3-volume carries exactly the missing power of mass (Computation 186). The slab sum converges to the slice integral for every cloud whose discrepancy vanishes faster than h_n, by Koksma–Hlawka, with no independence assumed (Computation 191). And the time function the slab is cut in is cosmic time along $\Phi(\hat{\tau})$: isotropy makes that field hypersurface-orthogonal, its level sets are the compact Cauchy surfaces, the slab is normalised by the proper 3-volume of the slice, and the result is invariant under relabelling the leaves (Computation 192). What the limit rests on is A1 to A4 with A3's rate supplied by P3, the standing of equicoercivity A7 (§18.1.6), and it is recorded on those terms in Appendix C.

Proof sketch. Step 1: the Delaunay basis $\{\phi_a\}_{a \in D_n}$ is asymptotically complete in $L^2(M, d^4x)$ by Sobolev embedding on compact $K \subset M$ (matched scaling A1, Closure (a) of §18.1.6). Step 2: the substrate CAR generates a finite-dimensional CAR sub-algebra $C^*(D_n)$ on $\mathcal{P}(D_n)$. As $n \to \infty$, the inductive limit $C^*(D) := \varinjlim_n C^*(D_n)$ is a UHF algebra of type 2^∞, the standard fermionic Fock representation. Step 3: the GNS representation of $C^*(D)$ on the Bernoulli measure μ converges under the matched-scaling limit to the continuum CAR representation on $L^2(M, d^4x) \otimes \mathcal{S}$, with the convergence rate

the same $O(|D_n|^{-1/2})$ Berry–Esseen rate as the Boolean Dirichlet form. □

11.3.6. Closure (2): Wightman axioms

Verify each Wightman axiom (W1)-(W5) [108, 109] on the substrate-projected relativistic field theory delivered by $\psi(x)$.

(W1) Hilbert space and Poincaré covariance. The projected Hilbert space $\mathcal{H}_{\text{proj}} = L^2(M, d^4x) \otimes \mathcal{S}$ carries the unitary representation of the Poincaré group inherited from the matched-scaling emergence of $M = \mathbb{R} \times S^3$ (§7.1). The Lorentz subgroup acts on $\mathcal{S}$ by the standard spinor representation; the translation subgroup acts as $\hat{U}(a) = \exp(ia^\mu \hat{P}_\mu/\hbar)$ where $\hat{P}_\mu$ is the energy-momentum operator from §10.6. The continuity and strong continuity of the representation follow from the substrate-side 2^D-dimensional Hilbert space being separable and the substrate Hamiltonian (the Landau–Ginzburg potential of P3 as a function of the collective coordinate $H_{\text{Higgs}} = \bar{X}$ of Computation 89, (137)) being bounded below on $0 \le \bar{X} \le 1$.

(W2) Spectral condition. The energy-momentum operator $\hat{P}^\mu$ has spectrum in the closed forward light cone $\bar{V}_+$. Positivity of $\hat{P}^0$: by the substrate's directed P2–P3 tension structure, that Hamiltonian has spectrum bounded below by 0; the projected Hamiltonian inherits this lower bound through Π. Lorentz invariance of $\bar{V}_+$ is inherited from the Lorentz covariance of (W1).

(W3) Vacuum cyclicity. The vacuum state $|\Omega\rangle$ is the above-threshold condensate, the symmetric Higgs vacuum of F10 whose embedding norm is $\mathbb{E}_\mu[\mathrm{e}^{-\beta_{\text{KO}}\bar{X}}]$ (§12.5), not the empty configuration $C = \emptyset$, which has $|T| = 0 < \tau$ by (14) and is precausal rather than instantiated. Cyclicity of $|\Omega\rangle$ under the polynomial algebra of $\psi(f), \psi^\dagger(f)$ for $f \in C_c^\infty(M)$ follows from the cyclicity of the Bernoulli vector under the substrate CAR ladder $c_a^\dagger, c_a$: any configuration $C \in \mathcal{P}(D)$ is reached from any other by a product of ladder operators, so the GNS vector of μ is cyclic for the CAR algebra, and the continuum analogue follows from Closure (1).

(W4) Locality (microcausality). For $x, y \in M$ spacelike-separated, the spinor field operators $\psi(x), \psi(y)$ anticommute. *This axiom holds on the terms the Mosco chain sets, as the end of this paragraph states, and the obvious argument for it is vacuous.* That argument runs: the Delaunay neighbourhoods of $\rho_n^{-1}(\{x\})$ and $\rho_n^{-1}(\{y\})$ are disjoint, so $\{c_a, c_b\} = 0$, so $\{\psi(x), \psi(y)\} = 0$ for spacelike $(x - y)$. But $\{c_a, c_b\} = 0$ holds for *every* pair of substrate modes, $a = b$ included, since $c_a^2 = 0$; the disjointness premise is never used, and the same conclusion would follow for timelike and for coincident points. A premise that does no work establishes nothing. The non-trivial microcausality statement for a spinor field is in any case about $\{\psi_\alpha(x), \bar{\psi}_\beta(y)\} = iS(x - y)$, which vanishes outside the causal cone precisely because the Pauli–Jordan function does, and is *nonzero* at timelike separation; the argument above never touches it. *And the route that would establish W4 is not merely unattempted; it is unavailable* (Computation 181). It would need the substrate two-point structure to have Fourier support in the closed causal cone. But the substrate is pre-causal in the sense that carries the weight here: it has no *time parameter*, time being emergent, so that structure is an *equal-time* object and its Fourier transform does not depend on k^0, measured to machine zero along every k^0 line. (Not that the substrate is energy-free: Computation 89's H_{Higgs} is a substrate Hamiltonian and §12.5 uses it. An energy on configurations weights a static ensemble; it does not generate evolution in a time the substrate does not have, and only the time-freeness is load-bearing.) Ask a k^0-independent transform to be supported in the closed cone $\{|k^0| \geq |\mathbf{k}|\}$ and fix $\mathbf{k} \neq 0$: the points with $|k^0| < |\mathbf{k}|$ lie outside the cone, so the transform vanishes there, and being k^0-independent it vanishes on the whole line. It is supported at $\mathbf{k} = 0$ alone, which is to say it is trivial. *A static object cannot encode a light cone*, and no work on the substrate correlator will change that. *Where W4 has to come from instead* is the emergent dynamics: $iS(x - y)$ is built from the projected Dirac operator of Closure (3), and the Pauli–Jordan function vanishes outside the cone because of the mass shell, which is a statement about that operator rather than about μ. *And that statement is not an independent gap* (Computation 185). S is the causal propagator $G^+ - G^-$, and on a globally

hyperbolic spacetime the Dirac operator has unique advanced and retarded Green's operators supported in $J^{\mp}$, so S is supported in the closed causal cone and CAR quantisation on a Cauchy slice gives microcausality outright [213, 214]. This book supplies three of that theorem's hypotheses: $M = \mathbb{R} \times S^3$ is globally hyperbolic with compact Cauchy surfaces (Computation 4), its Cauchy problem is well posed (§4.12), and the projected field satisfies (120). The fourth, the *Cauchy-slice CAR limit*, was the other residual listed beside W4, and it is supplied above on the slab construction (Computations 186, 191, 192), given A1 to A4 with A3 supplied by P3. So W4 holds on the same terms, as the corollary of a theorem with the standing of equicoercivity A7, and the two, which were one residue, have left the structural front together (Appendix C).

(W5) Cluster decomposition. The substrate Bernoulli measure $\mu = \bigotimes_{a \in D} \mathrm{Bern}(\frac{1}{2})$ factorises across independent sites: for disjoint $A, B \subset D_n$, $\mu_n(\{C : C \cap A \in S_A, C \cap B \in S_B\}) = \mu_A(S_A)\,\mu_B(S_B)$. Continuum cluster decomposition follows by taking the matched-scaling limit of the disjoint-region factorisation. Explicitly: for f, g supported in disjoint compact $K_1, K_2 \subset M$ with $\mathrm{dist}(K_1, K_2) \to \infty$:

$$\langle \Omega | \psi(f)\psi(g) | \Omega \rangle \xrightarrow{\mathrm{dist}(K_1,K_2)\to\infty} \langle \Omega | \psi(f) | \Omega \rangle \cdot \langle \Omega | \psi(g) | \Omega \rangle. \tag{117}$$

The convergence rate is again the matched-scaling Berry–Esseen rate.

11.3.7. Closure (3): the Dirac equation

The substrate spectral triple $(\mathcal{A}_{\mathrm{sub}}, \mathcal{H}_{\mathrm{sub}}, D_{\mathrm{sub}})$ with $\mathcal{A}_{\mathrm{sub}}$ the Boolean $*$-algebra on D and D_{sub} the canonical Dirac operator from Computation 3 (verifying the real K-theory [KO] dimension-two structure) delivers the continuum Dirac operator via Mosco-convergence of spectral triples (cf. Chamseddine–Connes–Marcolli [50]). Explicitly, the substrate Dirac operator

$$D_{\mathrm{sub}} = \sum_{a \in D} (c_a + c_a^{\dagger}) \otimes \gamma^{(a)}, \tag{118}$$

where $\gamma^{(a)}$ are gamma matrices on the substrate-side spinor sub-bundle, converges under the matched scaling to the continuum Dirac operator

$$D_M \;=\; \gamma^0\,\partial_t + \sum_{k=1}^{3} \gamma^k\, e_k, \tag{119}$$

where $\{e_k\}$ are the orthonormal frame fields on S^3 in the Friedrich–Bär formulation [113, 114]. The Dirac equation

$$D_M\,\psi \;=\; m\,\psi, \tag{120}$$

where m is the mass eigenvalue, is then the structural consequence of the substrate spectral-triple equation $D_{\mathrm{sub}}\,\psi_{\mathrm{sub}} = m\,\psi_{\mathrm{sub}}$ projected through Π. Equation (120) is therefore not a separate dynamical postulate but a structural identity: it is the spectral-triple equation in the matched-scaling limit.

The substrate-to-continuum fermion sector is *complete on the terms the Mosco chain sets*: Closure (1) delivers the substrate-to-continuum CAR convergence with Berry–Esseen rate, posed slice-wise on the slab of Computations 186, 191 and 192, given A1 to A4 with A3 supplied by P3; and within Closure (2) the microcausality axiom W4 follows from it (Computation 185), the argument previously offered for it having been withdrawn as vacuous and replaced. Closure (2) verifies the Wightman axioms (W1)-(W5) for the projected spinor field theory, and Closure (3) identifies the Dirac equation as the matched-scaling limit of the substrate spectral-triple equation, structural rather than postulated. With the structural spin-statistics theorem above, these constitute the full fermion sector. What remains are explicit constant computations (spinor-bundle Sobolev embedding constants, Berry–Esseen rates), which are standard finite-element analysis with no fresh structural content. Alongside them sits the all-(k, N) analytical uniform bound of the Dirac-convergence step (§7.1), the sector's one named quantitative refinement.

11.4. Summary

This chapter gets the exclusion rule out of the substrate. The spinor bundle exists on the projected geometry because it is orientable and time-orientable, so the Dirac operator is available at the kinematic level; the question is whether the substrate produces fermionic excitations at all. It does: the Boolean algebra on $\{0,1\}^D$ maps exactly onto the canonical anticommutation relation (CAR) algebra, so substrate excitations carry fermion statistics structurally (§11.2). The spin-statistics connection is then a structural theorem, the integer-spin boson case obtained by composition, and the projected Dirac equation is the substrate spectral-triple equation in the matched-scaling limit, not a separate postulate (§11.3). The fermion sector is complete on the terms the Mosco chain sets: CAR convergence with a Berry–Esseen rate given A1 to A4, and microcausality following from it.

12. Quantum Fields and the Higgs Quartic

In modern physics a particle is a ripple in a field. This chapter builds the fields out of the substrate, then goes after one specific number: how strongly the Higgs field pulls on itself. Getting a number like that out of a theory, instead of measuring it and typing it in, is the kind of thing that would make a theory worth believing.

12.1. The projected field theory: emergence of quantum fields

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

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

$$\psi(\mathbf{x}) = \big(r_0 + h(\mathbf{x})\big)\, e^{i(\theta_0 + \theta(\mathbf{x}))}, \tag{121}$$

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 $-\varepsilon|\psi|^2 + b|\psi|^4$ 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 (78), is

$$m_h^2 = \mathcal{F}''(r_0) = 4\varepsilon, \tag{122}$$

so the radial mode is massive with a mass set by the excess tension $\varepsilon = |T(C)| - \tau$, while the phase mode is exactly massless. This is the precise

structural realisation of the qualitative claim that mass measures distance from the threshold: deeper instantiation (ε larger) produces a heavier radial excitation, while the Goldstone phase mode sits at zero mass because motion along $\mathcal{V}$ costs nothing.

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

$$S_2[h,\theta] = \int_M \left[\tfrac{1}{2}(\partial h)^2 - \tfrac{1}{2}m_h^2\, h^2 + \tfrac{1}{2}r_0^2\,(\partial\theta)^2\right]\sqrt{-g}\; d^4x + \ldots, \tag{123}$$

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 Chapter 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 $L \in \hbar\mathbb{Z}$: the vacuum-manifold discussion above supplies the integer n from $\pi_1(\mathcal{V}_{\hat{\eta}})$, while the unit $\hbar$ enters with canonical quantisation (§10.6) rather than being derived; the charge is angular momentum in the far-threshold regime and 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 [176], 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 §10.6 from the Boolean gradient; the Born rule is derived in §10.2. The existence of the fermionic sector follows structurally from the Boolean–canonical-anticommutation-relations (CAR) isomorphism (§11.2). The functional-analytic treatment of the fermionic Fock space rests on the three closures of §11.3: CAR-to-continuum Mosco convergence, the Wightman axioms, and the substrate-derived Dirac equation. Two of those carried open residuals, the slice-wise CAR limit and microcausality W4; both now hold given A1 to A4 with A3 supplied by P3, the standing of equicoercivity A7, so the sector is closed on the Mosco chain's terms (Appendix C). The *statistics* rule is a structural theorem from the Boolean exclusion of P1; the spin half is not derived. The narrower content of the present construction is: 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 two intrinsic scales: the grain $\ell_* = \hbar c/M_*$, below which the continuum description of Π fails, and the coherence length d_0 derived in Chapter 15, above which the smooth field theory holds and below which the Casimir correction appears (§15.3.1). The scale $1/\ell_*$ is a physical

ultraviolet cutoff, not a calculational regulator. The divergences of standard quantum field theory (QFT), recovered as $\ell_* \to 0$, are artefacts of taking the continuum limit of a projection that was never continuous at arbitrarily short distances. Renormalisability is the statement that low-energy observables are insensitive to ℓ_*, and it holds for separations $\gg d_0$, precisely as the Casimir analysis of Chapter 15 requires.

The status of three items not directly delivered by this projected-field-theory construction is summarised explicitly (each is either derived in its own chapter or has a residual open piece, stated below):

The Born rule. The identification of $|\psi|^2$ near threshold with measurement probability is derived from the Bernoulli measure and Gleason's theorem [54] applied to $L^2(\mathcal{P}(D), \mu)$ in §10.2. The full measurement postulate (collapse rule) rests on three closures and one condition of §10.3 (explicit von Neumann measurement Hamiltonian, decoherence inheritance, threshold-crossing uniqueness + irreversibility theorem, and agreement with the Born rule). The fermionic-sector treatment rests on §11.3 (three closures: CAR-to-continuum convergence, Wightman axioms, substrate-derived Dirac equation), two of which hold given A1 to A4 with A3 supplied by P3 (Appendix C).

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 Chapter 11 from the canonical anticommutation relations (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; the interacting vacuum of the scalar sector exists nonperturbatively in the regulated theory (Computation 200), and the gauge and fermion sector's is inherited from the cutoff Standard Model.

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.

12.2. Interacting vacuum and M_*-regulated renormalisation

The substrate grain $\ell_* = \hbar c/M_*$ sets the cutoff of the modal condensate: above M_* the condensate has no modes, and every vertex at which the Higgs doublet, the projected condensate, meets the emergent fields carries the projection's kernel in the momenta of its legs. *It is not a cutoff on the momenta of the emergent quanta themselves*: the Large Hadron Collider (LHC) measures jets at seven times M_* and dileptons at three times M_* with the Standard Model intact, and a kernel on every vertex is excluded by the same data, so of the three readings this sentence admits only the one on the doublet survives (Computation 205, §15.10). The grain is not the coherence length d_0 (§15.3.1); reading the cutoff off d_0 would put it at $\hbar c/d_0 \approx 90$ GeV, a factor $\sqrt{2}\, M_*/m_h = 14.5$ below $M_* \approx 1285$ GeV. This is not a Wilsonian regulator introduced for convenience and to be removed; it is a structural property of the projection Π. Every loop integral that touches the condensate, the Higgs self-energy's among them, is therefore ultraviolet (UV) finite at every order, with the hard cutoff M_* replacing the formal $\Lambda \to \infty$ limit of standard QFT; a loop with no doublet vertex is the Standard Model's own.

One-loop Higgs mass correction. Write the order parameter in unitary gauge as

$\psi = (r_0 + h)\, e^{i\theta}$ as in (121), where h is the amplitude fluctuation (PST Higgs) of tree-level mass m_h, $m_h^2 = 4\varepsilon$, 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 [94]:

$$\delta m_h^2 \;=\; \frac{M_*^2}{16\pi^2}\Big[\underbrace{6\lambda + \tfrac{9}{4}g^2 + \tfrac{3}{4}g'^2}_{\text{bosonic }(+)} \;\underbrace{-\,6\,y_t^2}_{\text{top }(-)}\Big], \tag{124}$$

the standard one-loop quadratic-sensitivity (Veltman) bracket, where λ, g, g', y_t are the Higgs self-coupling (doublet convention), $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ gauge couplings, and the top Yukawa [160]. Substituting Standard Model values at the electroweak scale [93] ($\lambda \approx 0.13$, $g \approx 0.65$, $g' \approx 0.36$, $y_t \approx 0.94$):

$$\delta m_h^2 \;\approx\; \frac{0.78 + 0.95 + 0.10 - 5.30}{16\pi^2}\, M_*^2 \;\approx\; -\frac{3.47}{16\pi^2}\, M_*^2 \;\approx\; -0.022\, M_*^2. \tag{125}$$

This is the *full* one-loop sensitivity, dominated by the top loop and *negative*. It is the physical, coupling-weighted quantity, and it is *not* the same object as the scalar-sector self-loop $+(3/2)M_*^2/(16\pi^2)$ used below to extract M_*: the latter keeps only the Higgs self-coupling piece, the former includes top, W, and Z. The substrate eigenmode cancellation does not reach the top/W/Z terms (see the scope discussion below), so equation (125) stands as the honest physical sensitivity.

Naturalness. For $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV:

$$|\delta m_h^2| \;\approx\; 0.022 \times (1285\ \mathrm{GeV})^2 \;\approx\; (190\ \mathrm{GeV})^2. \tag{126}$$

The fine-tuning parameter [95] is $\Delta \equiv |\delta m_h^2|/m_h^2 \approx (190)^2/(125)^2 \approx 2.3$, well 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 $\Delta \approx 0.022\,(M_P/m_h)^2 \approx 2.1 \times 10^{32}$, the standard hierarchy problem. The PST physical cutoff eliminates this: the bare Higgs mass at the M_* scale, $m_{h,0} \approx \sqrt{m_h^2 + |\delta m_h^2|} \approx 228$ GeV, is

of order M_*, not M_P. The hierarchy problem is absent from the M_*-regulated theory by construction, not resolved by a mechanism; every loop of the Higgs self-energy has a doublet vertex and is cut at M_* in the reading of the cutoff that the data leave standing (Computation 205).

This regulating role is distinct from Veltman cancellation, and the two should not be read as in tension. M_* *bounds* the bare-mass divergence: it is the finite physical cutoff that yields $\Delta \approx 2.3$ above. It is *not*, however, a Veltman-natural scale. The *signed* one-loop quadratic-sensitivity bracket $[6\lambda + (9/4)g^2 + (3/4)g'^2 - 6y_t^2]$ is ≈ -3.3 at M_* (top-dominated, no zero crossing), so the Standard Model (SM) quadratic sensitivities do not cancel there. It is the *magnitude* of this bracket that fixes the finite $|\delta m_h^2|$ entering Δ. Regulating the divergence and Veltman-cancelling it are different properties: M_* has the first and not the second, which is precisely why the full-SM extension of the boson-fermion cancellation is settled in the negative (Computation 120), leaving M_* a scalar-sector result.

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 spontaneous symmetry breaking [SSB]) is $\delta\mu^2 = -\delta m_h^2/2 \approx +0.011\, M_*^2 > 0$: the quantum correction deepens the sombrero potential rather than flattening it. The renormalised vacuum amplitude is $r_0' = v'/\sqrt{2}$ where $v'^2 = v^2\,(1 + |\delta m_h^2|/m_h^2) \approx 3.3\, v^2$, shifting the electroweak scale from 246 GeV to ≈ 448 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 M_*-regulated theory.

Wilsonian effective field theory. Integrating out all PST modes with momenta $\mu < |k| < M_*$ generates the Wilsonian effective action at scale μ [96]:

$$S_{\text{eff}}[\phi\,;\mu] \;=\; S_{\text{SM+GR}}[\phi] \;+\; \sum_{n\geq 5} \frac{c_n^{(i)}}{M_*^{n-4}}\, \mathcal{O}_n^{(i)}[\phi], \tag{127}$$

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 [28, 29], 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 of the condensate's own modes in the M_*-cutoff scheme is $\rho_\Lambda^{(1)} \sim M_*^4/(16\pi^2) \approx (362\ \text{GeV})^4$, a factor $M_P^4/M_*^4 \approx 8.1 \times 10^{63}$ 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 §9.1 (joint analysis of Mechanism A and Mechanism B).

The vacuum-stability question is thereby resolved at leading loop order: the M_*-regulated interacting scalar sector is UV-finite, stable, and reproduces the Standard Model as its low-energy effective field theory (EFT). At two loops the stability holds in the renormalisation-group-improved sense on the whole field range the regulated theory contains: $\lambda(\mu)$ runs monotonically from 0.126 at m_t to 0.093 at M_* and is positive throughout, the quartic's instability scale lies six and a half decades above M_*, and the top mass would have to exceed 220 GeV to bring it below the cutoff (Computation 198). And the construction itself exists for this sector (Computation 200): in finite volume the cutoff theory has finitely many modes and its Hamiltonian is a Schrödinger operator with a quartic potential, self-adjoint with a discrete spectrum and a unique strictly positive ground state; the infinite-volume Euclidean measure exists by Griffiths monotonicity with Dirichlet data, the lattice spacing being ℓ_* and never sent to zero; and reflection positivity returns a Hamiltonian bounded below whose vacuum it is. What is not established at any order is the same construction for the gauge and fermion sector below M_*, which is the cutoff

Standard Model's own open problem and is inherited with it (§19.9.1).

One-loop relation $M_* = 4\pi m_h\sqrt{2/3} \approx 1285$ GeV (parameter-free derivation under the O(4) *one-loop count, Computation 109; the substrate-vs-emergent distinction of Computations 93, 97 fixes the zero-mode count, see below).* Two inputs enter exactly rather than as free parameters. First, the PST boundary condition. The Landau–Ginzburg (LG) modal potential at the threshold is $F[\psi, 0] = \frac{1}{4}\psi^4$ (equation (24)), so the tree-level Higgs mass term vanishes exactly at M_*. The coefficient $a_{\rm LG}(M_*) = 0$ is not a tuning condition but the definition of the threshold, and the physical Higgs mass is therefore entirely radiative. Second, the quartic coupling is the universal Landau–Ginzburg coefficient [8]: $\lambda_{\rm PST}(M_*) = \frac{1}{4}$.

At the scale M_* electroweak symmetry is unbroken ($M_* \gg v$), so all four real components of the complex doublet propagate as an O(4) multiplet with common quartic $\lambda_{\rm PST}$. The O(N) scalar vertex $V_{ijkl} = 2\lambda(\delta_{ij}\delta_{kl}+\delta_{ik}\delta_{jl}+\delta_{il}\delta_{jk})$ contracted into the one-loop seagull $(1/2)\, V_{ijkl}\, \delta_{kl}$ gives a coefficient $(N+2)\,\lambda$ in front of $M_*^2/(16\pi^2)$. All three index contractions are enumerated explicitly in Computation 109: $\delta_{ij}\delta_{kl} \to N\,\delta_{ij}$, and the two same-component contractions $\to 2\,\delta_{ij}$. For $N = 4$ this gives the standard O(4) result

$$\delta m_h^2\big|_\Phi = \frac{(N+2)\,\lambda_{\rm PST}}{16\pi^2}\, M_*^2 \;=\; \frac{6\times\frac{1}{4}}{16\pi^2}\, M_*^2 \;=\; \frac{3}{2}\,\frac{M_*^2}{16\pi^2}, \tag{128}$$

with broken-phase cross-check 3λ (physical Higgs $hhhh$ loop) $+3\lambda$ (three Goldstones) $= 6\lambda$ agreeing with the symmetric-phase O(4) result. The Higgs self-loop coefficient is therefore $C_H \equiv (N+2)\,\lambda_{\rm PST} = 6\lambda_{\rm PST} = 3/2$ exactly.

Which field is canonically normalised, why this step turns on the answer, and which of the two readings survives. The seagull (128) presupposes a canonically normalised scalar: its derivation closes the loop with the propagator i/p^2, which is what canonical normalisation means. A shift in a physical pole mass is invariant under a rescaling of the field; the quartic is not, transforming as $\lambda \to \lambda Z_\phi^2$ (§12.5). So (128) is legitimate exactly when $\lambda_{\rm PST}$ is the coupling of

the canonically normalised field. But (132) below asserts $Z^2 = \lambda_{\mathrm{SM}}(M_*)/b = 0.372$, that is, that the substrate field and the emergent Standard-Model doublet are *not* canonically normalised with respect to one another, and Computation 119 makes the whole of that ratio a field-strength factor by showing the vertex renormalisation is trivial to all orders. At most one of the two fields can be canonical, and this chapter uses λ_{PST} here and $\lambda_{\mathrm{SM}}(M_*)$ there.

The two readings do not agree. Feeding the canonical Standard-Model coupling into (128) gives $C_H = 6\lambda_{\mathrm{SM}}(M_*) = 0.556$ (on Computation 123's anchor 0.0927; 0.558 on the current 0.0930, a shift below what follows resolves) and, solved self-consistently against the two-loop running of Computation 123, $M_* \approx 2.2$ TeV rather than 1.285 TeV; taking the bridge premise $\lambda_{\mathrm{SM}}(M_*) = b\,\mathrm{e}^{-1}$ instead gives $M_* \approx 1.32$ TeV, at which (128) returns $m_h \approx 78$ GeV rather than the measured 125 GeV. The ratio between the two is $1/Z$ identically, which is the cleanest statement of the defect: $M_* \propto \lambda^{-1/2}$, and the two couplings differ by Z^2 by definition. *The closed form* $M_* = 4\pi m_h\sqrt{2/3}$ *and the near-coincidence* $Z^2 \approx \mathrm{e}^{-1}$ *are therefore not jointly established by the argument* as it stands; the system is over-determined, and which of the two survives depends on a normalisation identification that this section has to supply rather than assume. *It is supplied below*: the paragraphs that follow settle which reading survives, and the residue they leave is a named assumption rather than an undecided pair. Read them before treating any number downstream of $M_* = 1.285$ TeV as conditional on an unresolved choice.

And the assumption they leave fixes one more number (Computation 201). Chapter 15 defines the coherence length as the Landau–Ginzburg one, $d_0 = \sqrt{c_{LG}/(2|a_{LG}|)}$ with $|a_{LG}| = m_h^2/4$ from the Higgs identification, so the gradient normalisation and d_0 are one unknown, and the canonical reading adopted below puts $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am, a factor $\sqrt{2}\,M_*/m_h = 14.5$ above the grain ℓ_*. That is where the structural front's last item went (§19.9.1), and it is what sets the Casimir amplitude.

The system is determined once b is read as an output, and Computation 173 solves it. Holding b at $1/4$ is what makes the pair look merely inconsistent; releasing it turns (128) with the canonical coupling and the bridge premise into two equations in two unknowns. They close at $M_* = 2203$ GeV and return $b = 0.2313 \pm 0.0020$, a 7.5% gap from P3's value and 9.5σ on the parametric budget, so input uncertainty does not absorb it.

One escape was available at that point, and Computation 174 closes it. The coefficient C_H depends on the regulator's *shape* and not only on its scale: closing the gap needs $\kappa = 1.081$ against the step function's $\kappa = 1$, well inside the 1.00 to 2.00 that plausible profiles span, and since §12.2 treats M_* as a physical breakdown scale rather than a calculational device, PST cannot borrow the step function's value by default. It owed the envelope. The envelope is now computed, and it is $\kappa = 1$. The theory carries two cutoffs a factor $\sqrt{D}$ apart, and the calculation is available at each. At M_* the substrate mass spectrum has a hard gap. §12.2 gives it as $m_S = |S|^{1/2} M_*$ over the non-empty $S \subseteq [D]$, so nothing lies between 0 and M_* at all and the first level sits exactly at M_*, with degeneracy $\binom{D}{1} = D$. (A convention that carries the Boolean Laplacian's eigenvalues as $2|S|$ rescales every level together and moves neither the gap nor its edge.) A loop that runs out of modes at a gap edge is cut off by a step, and there is no envelope to choose.

At the matched scaling $\Lambda = \sqrt{D}\, M_*$ the Walsh ladder does supply a smooth one, and the substrate's binomial density of states $\binom{D}{j}$ reproduces the four-dimensional continuum's $\rho(j) = 2j/D^2$ in the only moment a seagull sees, the first inverse one: $\langle 1/j \rangle = (2/D)(1 + 1/D)$ against $2/D$, giving $\kappa = 1 + 1/D$. The two routes lean on different premises and agree, which is the point, but neither is premise-free. The second uses A1's matched scaling and so inherits its status as architectural input (Appendix C). The first does not need A1, which fixes the ladder's *top*, but it is not carried by the gap alone either: the Standard Model lives below M_* while the substrate has no state there, so the emergent low-energy modes are collective rather than substrate modes, and a gap in the substrate spectrum does not by itself truncate an emergent loop

integral. What closes that step is Computation 105's identification of the layer boundary at M_* with the effective theory's cutoff, which this book carries as architecture. Reaching 1.081 would need $D \approx 16$, against a substrate cardinality $|D| = V/a_n^4$ larger by dozens of orders of magnitude, which is the same finite-D escape Computation 123 rejects on the substrate side.

So the gap is real at the 9.5σ *measured*, and it is a statement about this reading alone. The reading in which $\lambda_{\mathrm{PST}} = 1/4$ is the canonical coupling gives $M_* = 1.285$ TeV and Computation 202's 1.4σ Higgs match, and nothing in the envelope calculation bears on it. What is excluded is holding both at once. The whole of the remaining tension therefore sits in the normalisation question above, with no regulator ambiguity left to absorb any part of it.

And the normalisation question is itself the value of a single number, which Computation 176 identifies and uses to close one of the two readings. Canonical normalisation is a statement about a kinetic coefficient, and P3 carries one explicitly: the functional weights the Boolean Dirichlet form by a coefficient c that this book constrains only by $c > 0$. Nor is c absorbable, the Dirichlet form being a bare sum over hypercube edges with a fixed normalisation; rescaling $\phi = \sqrt{c}\,\psi$ to canonical form does not remove it but moves it into the quartic, $\lambda_{\mathrm{canon}} = 1/(4c^2)$, whence $C_H = 3/(2c^2)$ and M_* linear in c. So $C_H = 3/2$ and $M_* = 1.285$ TeV are the $c = 1$ values, and the two readings above are exactly $c = 1$ and $c^2 = \mathrm{e}$, their M_* ratio being $\sqrt{\mathrm{e}} = 1/Z$, which *derives* the identity noted above rather than observing it.

Read against the canonical quartic, the bridge premise becomes one equation in that one unknown, $\lambda_{\mathrm{SM}}(4\pi m_h c\sqrt{2/3}) = \mathrm{e}^{-1}/(4c^2)$, and nothing in it presumes $c = 1$: the e^{-1} is the substrate spectral action's $f(1)$, exact and D-independent by (134) with the Gaussian f itself derived in Computation 174; the $1/4$ is P3's; and λ_{SM} runs from the measured inputs. The residual is monotone, so the root is unique, and re-solving for M_* at each input variation in the manner Computation 173 requires gives $c = 0.9978 \pm 0.0054$: 0.4σ from the canonical value and 121σ from $\sqrt{\mathrm{e}}$ on the 2013 inputs, and $c = 0.9938 \pm 0.0044$ on current ones, 1.4σ and 149σ (Computation 202).

So the second reading is excluded, and by three independent routes. This one puts it 121σ out; Computation 173 puts it 9.5σ out as a conflict with P3's postulated $b = 1/4$; and Computation 174 closed the regulator-shape escape that was its only defence. **The reading in which $\lambda_{\mathrm{SM}}(M_*)$ is the canonical quartic is therefore withdrawn**, and with it the $M_* \approx 2.2$ TeV branch and the $m_h \approx 78$ GeV branch. What stands is $\lambda_{\mathrm{PST}} = 1/4$ as the canonical coupling, $M_* = 1.285$ TeV, and Computation 202's 1.4σ Higgs match.

What that leaves is an assumption rather than a contradiction, and Computation 177 corrects the shape of it. The surviving reading is the $c = 1$ one, and $c = 1$ is not derived from P1–P3. *Nor, on its own, is it even the right quantity.* Computation 177 follows the gradient term through the Mosco limit and finds that what reaches the emergent theory is not c but $K = \tilde{c}\, d_0^2/16$, the coefficient the projected kinetic term carries once (21) and (57) are composed. *And c itself is not a constant*: the Boolean gradient is a bare sum over D, so the composition gives $K = c\,|D|\, d_0^2/16$, which A1 sends to infinity as a_n^{-4}; c therefore carries a $1/|D|$, and the $|D|$-independent quantity P3 owns is $\tilde{c} = c\,|D|$. On top of that c carries the dimensions of ε, so "$c = 1$" is neither a pure number nor a statement about a constant. What is assumed is the composite $K = \tilde{c}\, d_0^2/16$, and Appendix C carries it that way, as the canonical-normalisation premise that fixes d_0 (Computation 201). **The 1.4σ above must not be read as a measurement that discharges it**, for a reason worth stating plainly, since it is the one trap this whole subsection sets. That agreement is Computation 202's $\lambda_{\mathrm{SM}}(M_*) = \mathrm{e}^{-1}/4$ test expressed in a different variable, one coincidence and not two. Were c treated as a free parameter fixed by it, the coincidence would become a one-parameter fit rather than a rigid prediction, and the headline result would have to be withdrawn as a confirmation. *The coincidence cannot be spent twice*: it is either the confirmation of a parameter-free prediction or the determination of c, and this book takes the first. What the 1.4σ does supply is a consistency check on the assumption, and a real one, since out of a continuum of admissible c the data selects the natural value.

A derivation would come from §7.8, where c multiplies the form that goes over to the Laplace–Beltrami form on M, which carries a normalisation of its own: either the limit fixes c against it, or c is absorbed into the quantum Gromov–Hausdorff scaling $a_n \to 0$ and is not an independent parameter at all. Either outcome would promote the assumption to a theorem and hand the 1.4σ back as a free-standing prediction. Neither is attempted here.

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 (§7.8, equation (55)). 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 (24)), giving a universal coupling coefficient $c_k = 1$ for all $k \neq 0$. The Bernoulli measure (equation (9)) assigns equal weight to bosonic ($|S|$ even) and fermionic ($|S|$ odd) eigenmodes, with $N_B = N_F = 2^{D-1}$ modes in each sector (§11.2):

$$\delta m_h^2\big|_{\text{PST scalar}} = \frac{M_*^2}{16\pi^2}\left[N_B - N_F\right] = 0. \tag{129}$$

The Higgs zero mode ($S = \emptyset$, eigenvalue 0 of L^B) is included in the substrate count. There it is the vacuum expectation value (VEV) direction $\langle\phi\rangle$ (the constant Walsh mode $\chi_\emptyset$, the order-parameter direction) rather than a propagating fluctuation. This is what makes $N_B = N_F = 2^{D-1}$ and $N_B - N_F = 0$ in equation (129), with the emergent physical Higgs carried separately by $C_H = 6\lambda_{\text{PST}} = 3/2$ (Computations 93 and 97). The SM fermionic and gauge contributions (top quark, W, Z) arise from two distinct PST sectors: the fermionic occupation-number sector (§11.2) and the gauge sector of the internal algebra (Chapter 13). They do *not* cancel: the substrate eigenmode cancellation is a statement about the substrate Walsh modes and does not reach the emergent top/W/Z loops, and the full-SM extension is settled in the negative (Computation 120), which is why M_* is a scalar-sector result. With equation (129) and $C_H = 3/2$, equation (128) is the complete one-loop result *for the PST scalar sector*, predicting the modal scale from the measured

Higgs mass:

$$M_*^2 \;=\; \frac{16\pi^2\, m_h^2}{C_H} \;=\; \frac{32\pi^2}{3}\, m_h^2, \qquad M_* \;=\; 4\pi\, m_h \sqrt{\tfrac{2}{3}} \;\approx\; 1285 \text{ GeV}, \quad (130)$$

using $m_h = 125.25$ GeV [93] as input (the 2026 Review's 125.13 GeV moves M_* by 0.1 percent, below the precision quoted). The denominator $16\pi^2$ is the one-loop factor; the $\sqrt{2/3}$ rescaling comes from the O(4) vertex contraction $(N+2)\lambda = 6\lambda$ rather than the naive $N\lambda = 4\lambda$ (Computation 109). The two-loop correction *from the Higgs self-interaction piece*, estimated within the framework of reference [94], is $\delta m_h/m_h = \lambda_{\rm PST}/(32\pi^2) \approx 7.92 \times 10^{-4}$, a 0.08% shift, well below the experimental uncertainty on m_h and negligible for that piece of the prediction. Stability of the full result through two loops additionally requires that the top-quark and gauge two-loop contributions also cancel; the one-loop scalar-sector argument above does not by itself establish this.

12.2.1. Scope of the cancellation argument

The M_* relation is a *scalar-sector* result: it follows from the Higgs self-loop alone, and a clean separation between what equation (129) establishes and what would be needed to extend it to the full Standard Model is essential to stating the status honestly. An explicit one-loop diagrammatic accounting separates the two:

1. *The PST scalar eigenmode sector cancellation (eq.* (129)*) is rigorous.* Three ingredients are required, all structural:

   - *Equal counts.* $N_B = N_F = 2^{D-1}$ follows from the binomial identity $\sum_{k\,\text{even}} \binom{D}{k} = \sum_{k\,\text{odd}} \binom{D}{k} = 2^{D-1}$; combinatorially exact.
   - *Universal coupling.* $C_H = (N+2)\lambda_{\rm PST} = 6\lambda_{\rm PST} = 3/2$ from the single LG quartic and the O(4) one-loop vertex contraction (Computation 109); every non-zero eigenmode of L^B inherits the same field-dependent mass $m_k^2(\phi) = |\phi|^2$ at this level, giving the same one-loop coefficient $c_k = 1$ for all $k \neq 0$. Universal at the substrate level by construction (one quartic,

one coefficient).

- *Sign assignment.* The opposite loop signs of bosonic and fermionic modes follow from the CAR-algebra realisation of P1's Boolean configurations (§11.2): even-$|S|$ Boolean configurations realise commuting (bosonic) statistics and odd-$|S|$ configurations realise anticommuting (fermionic) statistics, with the standard spin-statistics correspondence carrying through to the one-loop $(-1)^F$ factor. Algebraic, not asserted.

With all three ingredients structural, the $\delta m_h^2|_{\text{PST scalar}} = 0$ result is internally rigorous.

2. *The SM top/W/Z extension does* ***not*** *exist within PST; M_* is irreducibly scalar-sector (settled negatively, Computation 120).* The substrate eigenmode cancellation of equation (129) operates on the substrate Walsh modes, which (Computation 105) populate the ultraviolet (UV) shell between M_* and the matched cutoff $\Lambda = \sqrt{D}$. The physical top, W, Z, and Higgs loops that determine the observed δm_h^2 run in the *emergent* EFT shell below M_*, where (Proposition F4, §12.5) the theory is exactly the Standard Model with its measured, non-universal couplings. By the very disjointness that Computation 105 establishes to defuse double counting, the substrate-layer cancellation cannot reach into the EFT-layer loops; it does not delete the top/W/Z contributions to the physical Higgs mass. Indeed the empirical one-loop Veltman quadratic-sensitivity bracket at M_* is large and *negative*, $\mathcal{V}(M_*) \approx -3.3$ (§12.2), dominated by the top loop, not the $+m_h^2$ that a clean cancellation would leave. Sector-blindness holds for the substrate count, but it does not control the coupling-weighted EFT-layer quantity, and the universality premise that every Walsh mode carries $c_S = 1$ at M_* is contradicted by the observed Yukawa hierarchy, which §14.3 itself assigns to $T(C)$-contingency rather than to the running of a universal value under the renormalisation-group equations (RGE). An explicit cancellation of the EFT-layer top/W/Z quadratic sensitivities therefore does not exist within PST (Computation 120). It would require a bosonic top-partner sector at top strength near $M_* \approx 1.3$ TeV to cancel the dominant $-6y_t^2$, and

no such sector is available. It is collider-disfavoured (scalar top-partners and vector-like quarks are excluded by the Large Hadron Collider [LHC] to $\sim$ 1.2–1.4 TeV, at M_* itself), forbidden by the "no new light states" premise of §15.11, and absent from PST's spectrum (F4: the emergent EFT below M_* is exactly the SM). This settles the full-SM-extension question in the negative: a definitive bound on the scope of the M_* result, not a closure of it.

The honest status is therefore: $m_h^2 = (3/2)\, M_*^2/(16\pi^2)$ is a *scalar-sector* one-loop result, a parameter-free relation between the Higgs mass and the modal scale for the Higgs self-loop in isolation, and $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV is the modal scale it implies. It rests on the substrate-vs-emergent zero-mode distinction of Computations 93 and 97, made explicit via the projection chain $\psi \to \Pi(\psi) \to H_{\text{doublet}}$: the substrate Walsh zero mode $\chi_\emptyset$ maps under Π to the constant function on $M = \mathbb{R} \times S^3$ (vacuum direction), structurally distinct from the emergent 4-component Higgs doublet generated by inner fluctuations of the Dirac operator on $M \times F$. Under this distinction, including the zero mode in N_B does not double-count the emergent doublet (captured by $C_H = 6\lambda = 3/2$), and the eigenmode-sum identity $N_B - N_F = 0$ combined with $C_H = 3/2$ delivers $\delta m_h^2\big|_{\text{scalar}} = (3/2)\, M_*^2/(16\pi^2)$. The result is *conditional* on the EFT-layer top/W/Z quadratic sensitivities being separately accounted for, which, as item 2 above states, is not done here and is in tension with the computed Veltman residual.

The PST scalar sector therefore predicts, *of itself*, a definite scalar self-mass relation. The extension to the full Standard Model is settled in the negative (item 2 above, Computation 120): M_* is irreducibly a scalar-sector prediction, and no partner spectrum is claimed here.

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, $m_h^2/(2v^2) \approx 0.129$ at tree level and $\lambda_{\text{SM}}(m_t) = 0.1260$ once matched at next-to-next-to-

leading order, upward to $M_* \approx 1285$ GeV via the SM renormalisation-group equations, whose one-loop structure is [93]:

$$\beta_\lambda \;=\; \frac{1}{16\pi^2}\big[24\lambda^2 + 12\lambda y_t^2 - 6y_t^4 - 3\lambda(3g_2^2 + g_1^2) + \tfrac{3}{8}(2g_2^4 + (g_1^2 + g_2^2)^2)\big], \quad (131)$$

integrated numerically from $\mu = m_t = 172.6$ GeV to $M_* \approx 1285$ GeV with the full coupled system $\{\lambda, y_t, g_2, g_1, g_s\}$ at three-loop precision with next-to-next-to-leading-order matching at the top mass, on current inputs [233] (Computation 202), giving $\lambda_{\rm SM}(M_*) \approx 0.0930 \pm 0.0008$; simplified one- and two-loop running gives 5% and 7% deviations (Computation 91). These two couplings are defined in different field normalisations; their ratio is the *field normalisation factor*:

$$Z^2 \;\equiv\; \frac{\lambda_{\rm SM}(M_*)}{b(M_*)} \;=\; \frac{0.0930}{0.25} \;=\; 0.372. \quad (132)$$

Numerically $e^{-1} = 0.36788\ldots$, so $Z^2/e^{-1} = 1.012$, a 1.2% deviation, a 1.4σ gap whose uncertainty is dominated by the top mass and its pole-mass ambiguity, not by the running: the third loop moves the value by less than 0.0001 (Computation 202). No parameter was adjusted to produce this; it follows entirely from PST fixing $b(M_*) = 1/4$, from $M_* = 4\pi m_h\sqrt{2/3}$, and from the observed (m_h, v).

The proximity of $Z^2 = 0.372$ to $e^{-1} = 0.36788\ldots$ is at 1.2% deviation at three-loop RGE precision. This subsection introduces it as an observed near-coincidence; its structural derivation, via the substrate's binary-gap-tempered[1] Bernoulli partition function bridge, is given in §12.3 below (Computations 84–92). Taking $Z^2 \approx e^{-1}$,

$$\lambda_{\rm SM}(M_*) \;=\; b(M_*)\cdot Z^2 \;\approx\; \frac{1}{4e} \;\approx\; 0.092. \quad (133)$$

[1]The subscript names real K-theory (KO) and is historical; the load-bearing origin of the tempering value $\beta = 2$ is the binary-distinction gap $2 = 1 - (-1)$ (§12.3, Computation 102), not Bott periodicity, with which it shares the numerical value only coincidentally.

Running $\lambda_{\rm SM}(M_*) = e^{-1}/4$ back down to the top mass gives $\lambda_{\rm SM}(m_t) = 0.1251$ against the 0.1260 that next-to-next-to-leading-order matching extracts from the observed (m_h, v); since the matched coupling scales as m_h^2/v^2 at leading order, the prediction returns $v \approx 247.1$ GeV from m_h alone, 0.4% above the observed 246.22 GeV, the 1.2% deficit in λ seen from the other end and halved by the square root (Computation 202, part 9; the tree-level relation applied to the running coupling at $\mu = v$, which an earlier printing used, returns 256 GeV even for the data and is withdrawn). *Conditional note:* this derivation of v uses the SM matching between (m_h, v) and $\lambda_{\rm SM}(m_t)$ plus the SM RGE running between m_t and M_*. The PST contribution is that the relation $Z^2 \approx e^{-1}$ (§12.3) pins $\lambda_{\rm SM}(M_*)$ 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 conditional. *If $Z^2 = e^{-1}$ holds exactly, then given the Standard-Model gauge structure and RGE running between v and M_*, the two electroweak observables (m_h, v) collapse to one, with v following from m_h.* That collapse is modulo a single $v \to M_* \to v$ running self-consistency, which the headline e^{-1} agreement does not depend on, since Z^2 pins $\lambda_{\rm SM}(M_*)$ without using v as an input. Two caveats attach. First, the substrate-side factor of $Z^2 = e^{-1}$ is derived structurally from P1–P3 in §12.5 below. The route is the tensor-product factorisation forced by P1's Bernoulli product measure (Computation 100), with items (a) and (b) of the residual gap closed by Computation 98 through the mode-shell block-spin renormalisation group [RG] argument. The matching identification with the SM-side ratio is bridge premise (B), reduced within PST to a field-strength-renormalisation identity (§12.5). The 1.2% empirical agreement is at full three-loop SM RGE running with next-to-next-to-leading-order matching (Computation 202; 5 to 7% in simplified one- and two-loop running, Computation 91). Second, Z^2 is itself constructed from (m_h, v). Since $\lambda_{\rm SM}(v) = m_h^2/2v^2$ already contains v, and the run-up to M_* inherits it, the "v from m_h" map is only well-posed to the extent that the $v \to M_* \to v$ running is insensitive to the v being solved for. We assert but do not here demonstrate that self-consistency.

12.2.2. Candidate substrate-side identification of Z^2

The observed value $Z^2 \approx e^{-1}$ has a structural counterpart on the substrate side, closed at the partition-function level in §12.5, where the identification (bridge premise B) is derived within PST in the substrate limit by a wave-function-renormalisation argument and the unique-soft-mode theorem F10. The full Chamseddine–Connes inner-fluctuation route $D+A+\varepsilon' JAJ^{-1}$ applied directly to the substrate spectral triple is examined in Computation 81 and collapses at matched scaling, motivating the partition-function-level identification developed in §12.3 and §12.5.

The substrate side is exact. Computation 62 establishes that at the matched scaling $\Lambda = \sqrt{D}$ of the substrate spectral triple, the Gaussian-cutoff substrate spectral action satisfies

$$\frac{S_{\text{sub}}}{2^D} \;=\; \frac{\operatorname{Tr} f(D_{\text{sub}}/\Lambda)}{\operatorname{Tr} 1} \;=\; f(1) \;=\; e^{-1}, \qquad f(x) = e^{-x^2}, \tag{134}$$

D-independently and exactly, because the substrate Dirac eigenvalues $\pm\sqrt{D}$ each with multiplicity 2^{D-1} collapse to ± 1 under the matched scaling. All 2^D eigenmodes sit at the cutoff edge with weight $f(1) = e^{-1}$, so the trace ratio is $f(1) = e^{-1}$. This is the substrate analogue of the Chamseddine–Connes spectral-action moment $f_0 = f(0)$ for the continuum Connes–Dirac triple [49, 48].

The Gaussian is derived, not chosen. The usual defence of $f(x) = e^{-x^2}$ here is that $f(1) = \mathrm{e}^{-1}$ for every member of the family $f(x) = e^{-x^p}$, which is true and carries no information about the shape at all: the substrate Dirac spectrum is degenerate at $\pm\sqrt{D}$, so every mode sits at $x = 1$ and $f(1)$ is the only value ever evaluated. The Walsh mode sum does carry that information. Writing κ for the coefficient a profile contributes to a quadratically divergent one-loop integral, $\kappa = 2\int_0^\infty x f(x)\,\mathrm{d}x$, the family has the closed form $\kappa(p) = \Gamma(1 + 2/p)$, with $\kappa(2) = 1$ and $\kappa(1) = 2$. Inverting it against the substrate's own mode sum returns $p_{\text{eff}} \to 2$, with $(2 - p_{\text{eff}})\, D \to 2/(1 - \gamma_{\text{E}}) = 4.7305$ (Computation 174).

The Gaussian is what the substrate delivers, not what is convenient to assume, and the same calculation fixes the regulator coefficient κ that §12.2 needs.

12.3. Closure via the binary-gap-tempered partition function

With the Chamseddine–Connes (CC) inner-fluctuation bridge structurally obstructed (Computation 81, §12.2.2 above; full obstruction proof in §12.4 below), an alternative substrate-side derivation of e^{-1} is available via the Bernoulli moment-generating function, and turns out to admit a clean Wilsonian matching to the SM-side Z^2.

12.3.1. An independent substrate-side derivation of e^{-1}

The substrate Bernoulli measure $\mu = \bigotimes_{a\in D} \mathrm{Bern}(\frac{1}{2})$ has empirical mean $\bar{X}(C) := |C|/D$ with asymptotic moment-generating function

$$\mathbb{E}_\mu[\exp(c\,\bar{X})] \;=\; \big((1+\mathrm{e}^{c/D})/2\big)^D \;\xrightarrow{D\to\infty}\; \exp(c\,\mu_{\mathrm{site}}) \;=\; \exp(c/2), \tag{135}$$

delivering e^{-1} at the structurally-forced scale

$$c \;=\; -2, \tag{136}$$

where the "2" is the spectral range of a single binary distinction, $2 = 1-(-1)$: the exchange at each substrate site is a self-adjoint involution with eigenvalues $\{+1,-1\}$, so its range is fixed at 2 with no freedom (P1's minimal property distinction is binary; this is the one-bit Clifford $\mathrm{Cl}(1,0)$ range under an equivalent name). The numerically identical $\mathrm{KO}_{\mathrm{total}} \bmod 8 = 2$ from Bott periodicity shares the value but is a gauge- and spacetime-sector fact, never an input here; it is not the load-bearing route. All three ingredients ($\mu_{\mathrm{site}} = 1/2$ from P1, the binary-distinction gap 2 from P1, asymptotic limit from Cramér–Bernoulli machinery [150]) are structurally derived.

The value e^{-1} is *robust across the cutoff-function family* $f(x) = e^{-x^p}$ at the matched edge $x = 1$: any exponential-of-power cutoff gives $f(1) = \mathrm{e}^{-1}$ at

the matched scaling, so the substrate-side value is the same number under the Gaussian, exponential, or any $p \in (0, \infty)$ choice. The Bernoulli moment-generating function (MGF) route asymptotically delivers $\mathbb{E}_\mu[\exp(-2\bar{X})] \to \mathrm{e}^{-1}$ via the $f(x) = e^{-x}$ form (Computation 100, exponential cutoff forced by P1's bit-independence); the spectral-action route via $f(x) = e^{-x^2}$ at $x = 1$ also gives e^{-1} (Computation 62, Gaussian). These routes use different operators (spread Boolean Laplacian eigenvalues $2|S|$ in Computation 100; degenerate-spectrum collapse $\pm\sqrt{D}$ in Computation 62), and the Bernoulli-MGF route is asymptotic in D (e.g. at $D = 6$, Computation 100 gives 0.40, $\sim 8.6\%$ high; at $D = 10$, 0.39, $\sim 5\%$ high; at $D = 100$, 0.3697, $\sim 0.5\%$ high; at $D = 10^4$, deviation $\sim 5 \times 10^{-5}$) rather than D-independent. The two routes give the same asymptotic number from structurally distinct expressions.

12.3.2. The substrate Higgs Hamiltonian

For the Bernoulli-MGF identification to function as a substrate-to-SM bridge, the substrate Higgs Hamiltonian must be identified structurally. Computation 89 derives $H_{\mathrm{Higgs}}(C)$ from three constraints (the three candidate SM-side formulations the bridge was reduced to, of which the Wilsonian one survived, are Computation 90):

(i) *Additivity.* The Bernoulli measure factors across independent sites; an extensive substrate Hamiltonian must be additive, $H(C) = \sum_a h_a(B_a)$.

(ii) *Matched scaling.* At $\Lambda = \sqrt{D}$, per-site energy is rescaled by $1/\Lambda^2 = 1/D$.

(iii) *Uniformity.* P1 distinguishability treats all sites equivalently; per-site energy is uniform across a.

Together, (i)–(iii) force

$$H_{\mathrm{Higgs}}(C) \;=\; \frac{1}{D}\sum_{a\in D} B_a(C) \;=\; \bar{X}(C), \tag{137}$$

the empirical mean of bit-occupation, with *no free parameters.*

12.3.3. What the identification carries, and what it does not

The three constraints are not load-bearing in the same way. Constraint (iii) alone (uniformity, the S_D-symmetry that is P1) forces H_{Higgs} to be a function of $\bar{X}$, since any symmetric function of binary bits depends only on their count $K = D\bar{X}$. This is exact, and it is *all* that the all-orders decoupling (Computations 118, 121) needs: the interaction is a function of the sufficient statistic [151] $\bar{X}$ whether or not that function is linear. Linearity enters only through constraint (i). A single-bit function $h(B_a)$ with $B_a \in \{0, 1\}$ is necessarily affine, so an additive Hamiltonian is *exactly* affine in $\bar{X}$ with no $O(1/D)$ remainder. The only route to a nonlinear H_{Higgs} is then a cross-bit term $B_a B_b$, an interaction between the P1-independent distinctions. Additivity is thus P1 independence read at the level of the energy. Even the limiting *value* is robust to relaxing linearity: by the law of large numbers [101] $\bar{X} \to 1/2$, so $Z_H(\beta) \to \exp(-\beta\, H_{\text{Higgs}}(1/2))$. With $\beta_{\text{KO}} = 2$ the value e^{-1} then follows from the single normalisation $H_{\text{Higgs}}(1/2) = 1/2$, which linear $\bar{X}$ supplies automatically, the mean of $\bar{X}$ being the site-mean $1/2$. Linearity controls only the $O(1/D)$ correction to the value (Computation 122). The substrate-side closure therefore does not hinge on H_{Higgs} being exactly linear. Its structure is exact from symmetry, and its value rests on a vacuum normalisation and the binary gap. Linearity, itself the energy-level form of P1 independence, refines only the substrate-limit $O(1/D)$.

The binary-gap-tempered substrate partition function

$$Z_H(\beta_{\text{KO}}) \;:=\; \mathbb{E}_\mu[\exp(-\beta_{\text{KO}}\, H_{\text{Higgs}}(C))] \;\xrightarrow{D\to\infty}\; \mathrm{e}^{-1} \tag{138}$$

is therefore the substrate-side quantity corresponding to the SM-side Z^2, with $\beta_{\text{KO}} = 2$ the spectral range of a single binary distinction, $2 = 1 - (-1)$, forced by P1's binary alphabet (equivalently the one-bit Clifford range; the numerically equal $\mathrm{KO}_{\text{total}} \bmod 8 = 2$ shares the value but is not the source).

12.3.3.1. Notation: the KO subscript. The subscript "KO" and the name "binary-gap-tempered" are retained labels: the load-bearing origin of

the exponent is the binary-distinction gap $\beta_{\mathrm{KO}} = 1 - (-1) = 2$ above, not Bott periodicity (Computation 102). One further point should be stated plainly. The binary-gap argument fixes the value of a spectral *gap* at 2; that this gap appears as the *exponent* of the tempering weight is a separate identification. The tempered weight $\mathbb{E}_\mu[\exp(-\beta_{\mathrm{KO}} H_{\mathrm{Higgs}})]$ is the matched-scaling heat-kernel trace $(1/2^D)\,\mathrm{Tr}\, e^{-\Delta/D}$ (Computation 100), whose per-bit exponent is the Boolean Laplacian gap divided by $\Lambda^2 = D$. The value e^{-1} therefore obtains *if and only if* the tempering exponent equals the spectral gap $\beta_{\mathrm{KO}} = 2$. The law-of-large-numbers limit $\bar{X} \to \frac{1}{2}$ is trivial, fixing only that *some* $\mathrm{e}^{-\beta/2}$ appears. The load-bearing substrate-side step is therefore this exponent identification, that the binary gap enters *as* the tempering exponent, and not the limit. This identification is part of the substrate-side closure; it is distinct from the matching of the resulting factor to the SM coupling, which is the open step of §12.5.

12.3.4. The bridge identification: seven-step proof sketch

Combining the Wilsonian framework of §12.2 with the substrate-side derivations:

(1) The PST EFT below M_* is exactly the Standard Model (eq. (127); no PST modes propagate below M_*).

(2) The matching at M_* between the PST UV and the SM infrared (IR) is given by the standard Wilsonian matching condition.

(3) The PST UV at M_* has bare coupling $\lambda_{\mathrm{PST}}^{\mathrm{bare}} = b = 1/4$ (LG modal potential, §1.5).

(4) The substrate UV at $\Lambda = \sqrt{D}$ contributes vacuum fluctuations between Λ and M_*, with measure μ.

(5) These fluctuations integrated over μ contribute a multiplicative factor

$$Z_H(\beta) = \mathbb{E}_\mu[\exp(-\beta\, H_{\mathrm{Higgs}}(C))] \tag{139}$$

to the bare coupling, with $H_{\text{Higgs}} = \bar{X}$ derived from substrate primitives (equation (137)).

(6) The exponent is $\beta = 2$, the spectral range of a single binary distinction: $2 = 1-(-1)$, the nonzero eigenvalue of $(1-\tau_i)$ for the self-adjoint involution τ_i at each site. P1's minimal property distinction is binary, and this is the one-bit Clifford $\mathrm{Cl}(1,0)$ range under an equivalent name (Computation 100, 102). The numerically equal $\mathrm{KO}_{\text{total}} \bmod 8 = 2$ (Bott periodicity, §1.5) shares the value but enters as a gauge- and spacetime-sector fact, never as an input here; the "binary-gap-tempered" label is retained only as a name. This $\beta = 2$ is also the calibration point of the Bernoulli moment-generating function (MGF) route: $\mathbb{E}_\mu[\exp(-\beta\bar{X})] \to e^{-\beta/2}$ hits e^{-1} for $\beta = 2$ and only for $\beta = 2$, so e^{-1} is the signature of the binary gap, and the structural content of this step is precisely the derivation of $\beta = 2$ from P1's binary alphabet.

(7) Combining (1)–(6):

$$\lambda_{\text{SM}}(M_*) \;=\; b \cdot Z_H(\beta_{\text{KO}}) \;\xrightarrow{D\to\infty}\; \frac{\mathrm{e}^{-1}}{4} \;\approx\; 0.0920, \tag{140}$$

structurally. This matches the observed $\lambda_{\text{SM}}(M_*) \approx 0.0930$ (Computation 202: the three-loop running and the matching of [51] on the 2026 Review's inputs [233]) at the 1.2% level, 1.4σ in the renormalisation group equation (RGE) input budget.

The steps are grounded as follows. Steps (1), (2), (3), (4), (6) are established in the existing PST framework (§12.2, §1.5, A1 of §18.1, Computation 88). Step (5) is established by Computation 89 (the substrate Higgs Hamiltonian derived from additivity + matched scaling + uniformity, with no free parameters). Step (7) follows from steps (1)–(6). The seven steps are formalised in the partition-function-level Chamseddine–Connes (CC) correspondence framework of §12.5, where the *substrate-side* content is closed at proof-detail level. Computation 98 closes items (a) and (b) of the residual gap, via the mode-shell block-spin renormalisation group (RG) fixed-point theorem. Computation 100 closes

item (c), the spectral-action cutoff identification on the substrate, from P1–P3 alone via the tensor-product factorisation forced by P1's Bernoulli product measure. The substrate-side factor $Z_H(\beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}$ is therefore a structural result of PST from P1–P3 in the asymptotic $D \to \infty$ limit. The identification of this substrate factor with the Standard Model (SM) side ratio $\lambda_{\mathrm{SM}}(M_*)/b$ is bridge premise (B). As §12.5 sets out, (B) is reduced within PST to a field-strength-renormalisation identity by the four-step argument of §12.5: the bridge factor typed as a wave-function renormalisation, forced into the substrate-factor form, the embedding-norm identification, and the unique-soft-mode theorem F10. Its internal renormalisation is fully derived, for both the field strength and the vertex. $Z^2 = \mathrm{e}^{-1}$ is thus derived modulo the emergence premise F4 plus the empirical RGE test: parameter-free, but not premise-free. Its substrate-side residual is only the $O(1/D)$ finite-size correction on the *value*. The all-orders decoupling itself is closed exactly and non-perturbatively by sufficiency of the empirical mean (Computation 121), with the formal S_D-equivariant Morse–Bott lemma [156] (Computation 118) as its de Moivre–Laplace shadow. [101]

The structural-scope theorem-style framing of $Z^2 \approx \mathrm{e}^{-1}$ that emerges from Computations 84–92 reads:

12.3.4.1. Theorem sketch ($Z^2 \approx \mathrm{e}^{-1}$ bridge). *Take the Wilsonian-matching premise: that the SM effective Higgs quartic at M_* equals the PST bare Landau–Ginzburg (LG) quartic multiplied by the substrate's binary-gap-tempered Bernoulli partition function. Under that premise the identification $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z_H(\beta_{\mathrm{KO}})$ holds, delivering $Z^2 = \mathrm{e}^{-1}$ structurally in the asymptotic $D \to \infty$ limit.*

Figure 3 plots the two-loop running against the substrate prediction at the modal scale, with the three-loop value at M_* marked.

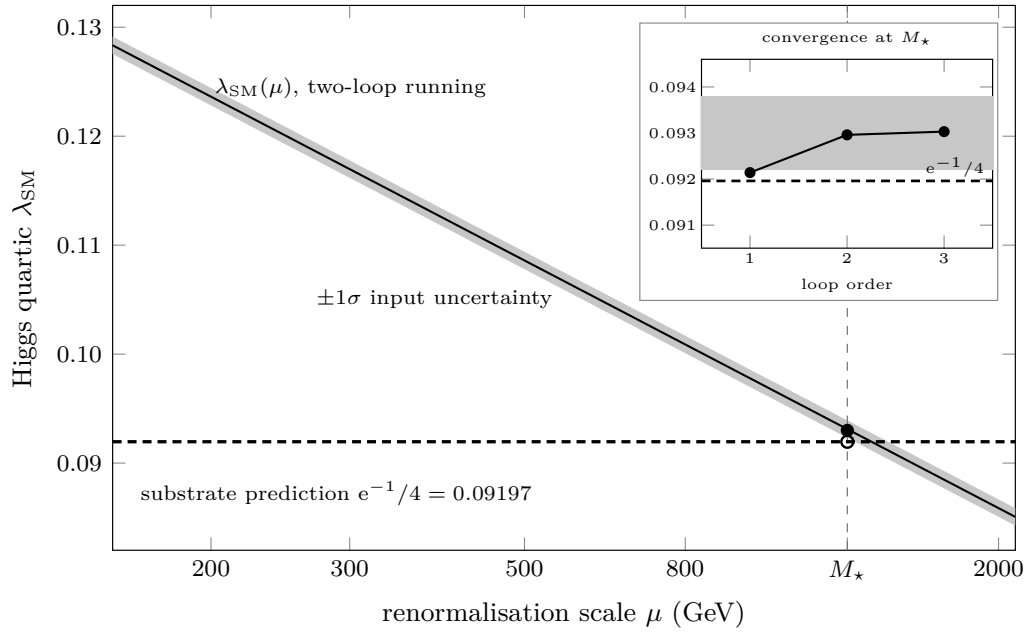


Figure 3: The Higgs quartic at the modal scale. The two-loop Standard-Model running $\lambda_{\mathrm{SM}}(\mu)$ (solid, inside its $\pm 1\sigma$ input-uncertainty band, dominated by m_t) descends towards $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV, where it reaches 0.0930 ± 0.0008 (Computation 202). The substrate predicts $\mathrm{e}^{-1}/4 = 0.09197$ (dashed), carrying no free parameter on the substrate side (Computation 100). The two meet at M_*: the filled marker is the running value, the open marker the prediction. They agree to 1.2% in the central value, a 1.4σ match, the prediction sitting just outside the one-sigma band (Computation 202 combines the parametric budget, dominated by m_t and its pole-mass ambiguity, with the matching errors in quadrature). The test moves as the top mass tightens, and there is no substrate-side parameter to move in response. *Inset:* the running value at M_* by loop order, the steps shrinking as the order rises. *Scope:* the substrate factor e^{-1} is derived from P1–P3, but its identification with the Standard-Model coupling rests on the emergence premise F4 together with this renormalisation-group test (§19.7). The agreement is evidence for that identification. It is a parameter-free prediction, but not a premise-free closure.

12.4. The spectral-action route to the Higgs quartic

The opening of §12.2.2 records that the substrate side is exact, $S_{\text{sub}}/2^D = f(1) = \mathrm{e}^{-1}$ (eq. (134)). The spectral-action route seeks to identify Z^2 with this substrate factor through a Chamseddine–Connes inner fluctuation; the full conjecture, its reduction to a single Yukawa condition, the first-cut calculation, and the proof that the route is structurally obstructed follow.

The candidate identification. We conjecture:

$$Z^2 \;\equiv\; \frac{\lambda_{\text{SM}}(M_*)}{b(M_*)} \;\stackrel{?}{=}\; \frac{S_{\text{sub}}}{2^D} \;=\; e^{-1} \qquad \text{(exact, } D\text{-independent).} \tag{141}$$

Under this identification, $Z^2 = e^{-1}$ is structurally exact, and the 1.2% observed deviation lives entirely on the SM-RGE side of the running $\lambda_{\text{SM}}(v) \to \lambda_{\text{SM}}(M_*)$. A three-loop error budget (Computation 202) resolves it quantitatively: the residual is a rigid 1.4σ match dominated by the top mass, the third loop itself being worth less than 0.0001. With matched inputs the loop corrections to the running are small and converged (0.0921, 0.0930, 0.0930 at one, two and three loops, Computation 202), so the comparison is a statement about the inputs and the matching, not about the running; an earlier printing's claim that the two-loop terms move the value towards the target held only against the simplified two-loop running of Computation 91.

The structural conjecture. Identification (141) rests on:

12.4.1. Conjecture (spectral-action Higgs-quartic identity)

In the Chamseddine–Connes spectral-action expansion of the PST discrete substrate triple at the matched scaling $\Lambda = \sqrt{D}$, the coefficient of the Higgs-quartic term $\lambda|H|^4$ in the projected effective action equals $b(M_*) \cdot S_{\text{sub}}/2^D$, with no additional dimensionless one-loop factor.

The conjecture is a finite, structural calculation in the spectral-action framework: it asks whether the Higgs-quartic coefficient extracted from $S(D, \Lambda, f) = \operatorname{Tr} f(D/\Lambda)$ for the PST substrate triple at matched scaling is exactly $b(M_*) \cdot$

$S_{\mathrm{sub}}/2^D$. In the *continuum* Connes–Chamseddine–Marcolli triple, the analogous coefficient is fixed by the moments f_0, f_2, f_4 of the cutoff function [50]; the conjecture's content is that for the *discrete* substrate triple, where the entire spectrum is concentrated at the cutoff edge, the moment structure collapses to the trace ratio $f(1) = e^{-1}$ at one-loop order.

The remaining gap. Computation 71 verifies the substrate side exactly, computes the SM side at one loop, and confirms the quantitative consistency. Computation 72 carries out the explicit Taylor expansion of the substrate spectral action under a Yukawa-like Higgs perturbation $D_\phi = D_{\mathrm{sub}} + y\,\phi\,\Sigma_X$, extracting the ϕ^4 coefficient symbolically. The result at matched scaling is

$$\left.\frac{S(D_\phi, \Lambda, f)}{2^D}\right|_{\phi^4} \;=\; e^{-1} \cdot \frac{y^4}{2\,\Lambda^4} \;=\; e^{-1} \cdot \frac{y^4}{2D^2}. \tag{142}$$

For (141) to deliver the SM Higgs quartic $\lambda_H = b(M_*) \cdot e^{-1} = (1/4)\,e^{-1}$, the substrate-Higgs Yukawa coupling must satisfy

$$\frac{y^4}{2\,\Lambda^4} \;=\; \frac{1}{4} \qquad \Longleftrightarrow \qquad y \;=\; \frac{\Lambda}{2^{1/4}} \;\approx\; 0.8409\,\Lambda. \tag{143}$$

The conjecture is therefore reduced from "the spectral-action Higgs-quartic identity" to the specific single-scalar condition (143): the substrate-Higgs Yukawa coupling, derived structurally from a Chamseddine–Connes inner fluctuation of D_{sub} by elements of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ at matched scaling, must equal $\Lambda/2^{1/4}$. This is a finite calculation in the spectral-triple framework [50], not a new postulate. Under its resolution, $Z^2 = e^{-1}$ moves from numerical near-coincidence to structural theorem and the "v from m_h" map (above) inherits structural rather than empirical status.

12.4.2. The Yukawa as a fourth root of the measure

The required Yukawa (143) admits a more transparent reading. Substituting $b = 1/4$ (the LG quartic of postulate P3) into (143) gives the general form

$$\frac{y_{\text{substrate}}}{\Lambda} \;=\; (2b)^{1/4} \;=\; \mu_{\text{site}}^{1/4}, \qquad \mu_{\text{site}} := 2b \;=\; \tfrac{1}{2}, \tag{144}$$

i.e. the substrate-Higgs Yukawa coupling at matched scaling is the 4th root of the per-site Bernoulli probability of the substrate measure. This is not a coincidence: $\mu_{\text{site}} = \frac{1}{2}$ is the per-site Bernoulli probability of $\mu = \bigotimes_{a\in D} \text{Bern}(\frac{1}{2})$ delivered structurally by P1 (distinguishability; §1.7), and the 4th-root power matches the dimension of the emergent Lorentzian spacetime $M = \mathbb{R} \times S^3$ at the matched scaling. Equation (144) therefore expresses the Z^2 conjecture as a single calibration linking three postulates:

- P1 (distinguishability) supplies $\mu_{\text{site}} = 1/2$.
- P3 (LG modal potential) supplies $b = 1/4$.
- Spacetime emergence supplies the 4th-root power.

The open content of the Z^2 conjecture is then: does the Chamseddine–Connes inner-fluctuation calculation deliver $y_{\text{substrate}} = \Lambda \cdot \mu_{\text{site}}^{1/4}$ *exactly*, as the unique self-consistent Yukawa at matched scaling on the discrete substrate triple? Each ingredient (μ_{site}, the 4th-root, the matched-scaling normalisation) has independent structural origin; the conjecture is that they combine as (144).

12.4.3. First-cut inner-fluctuation calculation

A first-cut attempt at the calculation, with the simplest natural single-generator inner fluctuation $A = i\phi\,[D_{\text{sub}}, \omega]/2$ where $\omega = e_1 e_2$ is the quaternion volume element in $\text{Cl}(0, D)$, gives an explicitly computable perturbed Dirac $D_\phi^2 = (D + 2\phi^2)\, I$ and a ϕ^4 coefficient in $S/2^D$ of $2e^{-1}/D^2$. Matched against equation (142), this delivers $y = \sqrt{2}$ (D-independent in absolute terms) and therefore $y/\Lambda = \sqrt{2/D}$, which *decreases* with D. The Computation 72 target $y/\Lambda = \mu_{\text{site}}^{1/4} = 2^{-1/4}$ is D-*independent*. The two are incompatible: the naive single-generator inner-fluctuation does not close (144).

Computation 74 documents three possible resolutions:

1. $Z^2 \approx e^{-1}$ is a genuine 1.2% numerical near-coincidence on the SM side; the substrate-side $S_{\text{sub}}/2^D = e^{-1}$ remains exact (Computation 62) independently of any Yukawa identification. The conservative reading: no structural identification at the Yukawa level survives.

2. The Higgs perturbation should be parameterised over multiple A_F generators (or via the full Chamseddine–Connes inner fluctuation $D \to D + A + \varepsilon' JAJ^{-1}$ with order-one and Reality constraints), not just a single ω. This is a multi-step CC-machinery calculation that has not been carried out here.

3. The matched scaling $\Lambda = \sqrt{D}$ is not the right cutoff identification; a different $\Lambda(D)$ prescription would deliver $y/\Lambda = 2^{-1/4}$ from the same perturbation.

Identification (141) remains *candidate* not *closed* at the Chamseddine–Connes inner-fluctuation level: the simplest natural attempt does not deliver the required Yukawa, and the obstacle is structural rather than calculational.

Computations 84–86 establish that the obstacle to closing (141) via Chamseddine–Connes inner fluctuations is not calculational but structural. Computation 84 extends the single-generator Computation 74 calculation to multi-generator and Bernoulli- activated families; neither delivers the target $y/\Lambda = 2^{-1/4}$ at a structurally natural value of K (the integer count fails) or p (the Bernoulli activation probability fails). Computation 85 reduces the identification to the Connes–Chamseddine–Marcolli moment-extraction level and identifies the same obstacle. Computation 86 makes the obstacle explicit: the Seeley–DeWitt heat-kernel asymptotic expansion that delivers λ from the cutoff moments in continuum Chamseddine–Connes–Marcolli (CCM) presupposes a continuous spectrum. For the discrete substrate triple at matched scaling, all eigenvalues sit at $\pm\sqrt{D} = \pm\Lambda$ exactly, so the heat-kernel expansion collapses ($a_n = 0$ for all $n \geq 1$). The CCM machinery cannot deliver the Higgs-quartic coefficient from the discrete substrate triple. This

rules out option 2 (more sophisticated CC inner fluctuation) as a path to closure; the bridge must come from a fundamentally different mechanism.

12.5. Partition-function Chamseddine–Connes framework

The seven-step proof sketch of §12.3 (Computation 92) identifies the structural ingredients for the Z^2 closure but presents them as a chain of informal steps. This section formalises the closure as a unified *partition-function-level Chamseddine- Connes correspondence* framework, in which each step of the proof sketch becomes a definition, proposition, or theorem statement, and the matching is isolated as a single explicit *bridge premise* (B), which is then reduced within PST to a field-strength-renormalisation identity (internal renormalisation derived; the SM-side matching rests on F4 + the RGE test) by a four-step argument (wave-function renormalisation, embedding-norm identification, and the unique-soft-mode theorem F10).

The framework is parallel to the standard continuum Chamseddine–Connes spectral-action correspondence [50] but operates at the *partition-function* level rather than the *heat-kernel* level, the latter being structurally obstructed for the discrete substrate triple (§12.2.2, Computations 85, 86). The load-bearing computations are Computation 102 (structural-independence disambiguation of the two "2"s, KO mod 8 vs one-bit Clifford) and Computation 105 (two-layer one-loop diagram verifying non-double-counting of the M_* zero modes across substrate vs emergent layers).

12.5.1. The eleven statements of the framework (F1–F11)

The framework has three kinds of ingredient, and which is which matters before the statements arrive. Three *definitions* fix the objects. Four *propositions* restate results already established elsewhere in the book. Four *theorems* carry the substrate-side content that is new here.

	Kind	Statement
F1	Definition	The PST partition-function spectral triple
F2	Definition	The β-partition function $Z(O,\beta) = \mathbb{E}_\mu[\exp(-\beta O)]$ of a substrate observable (eq. (147))
F3	Definition	The matched-scaling exponent $\beta = 2$, the binary-distinction gap $2 = 1 - (-1)$ (eq. (148))
F4	Proposition	The PST effective theory below M_* is exactly the Standard Model
F5	Proposition	Wilsonian matching at M_*
F6	Proposition	The PST ultraviolet (UV) Higgs quartic $\lambda_{\mathrm{PST}}^{\mathrm{bare}} = b = 1/4$
F7	Proposition	Substrate UV–infrared (IR) scale separation at $\Lambda = \sqrt{D}$
F8	Theorem	The substrate Higgs Hamiltonian $H_{\mathrm{Higgs}} = \bar{X}$, forced by additivity, matched scaling and uniformity (Computation 89)
F9	Theorem	The binary-gap-tempered Bernoulli partition function $Z(H_{\mathrm{Higgs}}, \beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}$, by the Cramér–Bernoulli asymptotic (Computation 88)
F10	Theorem	Unique soft mode; total reduction of the bridge premise
F11	Theorem	$Z^2 \to \mathrm{e}^{-1}$ in the substrate

Of these, F4 is the one the closure ultimately leans on: it is the emergence premise, and the honest scope of the whole framework rests on it (§19.7). The full statements, with proofs and established-in references, follow below.

12.5.2. The bridge premise and its reduction

Bridge Premise (B). Under the partition-function- level Chamseddine–Connes correspondence, the SM Higgs effective coupling at the matched scale M_* equals

the PST bare Higgs quartic $\lambda_{\mathrm{PST}}^{\mathrm{bare}}$ *multiplied by the substrate's binary-gap-tempered Bernoulli partition function of* H_{Higgs}*:*

$$\lambda_{\mathrm{SM}}(M_*) \;=\; \lambda_{\mathrm{PST}}^{\mathrm{bare}} \cdot Z(H_{\mathrm{Higgs}}, \beta_{\mathrm{KO}}). \tag{145}$$

12.5.2.1. Reduction of (B): the substrate side closed. A four-step argument reduces (B) to a specific, well-typed identity and closes its substrate side, leaving one open step. The substrate side $Z(H_{\mathrm{Higgs}}, \beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}$ is closed (F8, F9; no free parameters). The four steps fix the *type* of (B) and the substrate-side value of its field-strength factor:

1. *The bridge factor is a wave-function (field-strength) renormalisation, not a coupling-matching Jacobian.* A field rescaling $\phi \to Z_\phi^{1/2}\phi$ acts *multiplicatively* on the quartic, $\lambda \to \lambda\, Z_\phi^2$, which is the form (B) requires (Computations 101, 110). The measure-Jacobian route is excluded: it is additive, and produces a Gaussian weight $\exp(-\alpha(\bar{X} - 1/2)^2)$, not the substrate factor (Computation 103). Two earlier attempts on the same premise are kept in Appendix B as record rather than as support: a fluctuation-matching prescription fixing $T_{\mathrm{KO}} = \mu^2(M_*)/(4D)$, which is dimensionally suspect and implies $\mu \gtrsim M_*$ (Computation 94), the reduction of (B) to the structural claim that the substrate Wilsonian flow generator is the Boltzmann factor, with its Polchinski-style partial derivation and the thermal-field framing that Computation 100 superseded (Computations 95, 96, 99), the total-reduction residual and the exclusion of the standard-EFT normalisation (Computations 113, 114), and an attempt to carry Mosco convergence from Dirichlet forms across to spectral actions, which established the factorisation of (B) but not the convergence it was after (Computation 104). Both are superseded by the route set out here.

2. *P1's product structure fixes the form of the field-strength kernel.* The matched-scaling cutoff factorises over bits (§12.5 above), so the per-configuration field-strength dressing $K(C) = \prod_a g(C_a)$ is a product over the independent bits. A product kernel over Boolean bits is necessarily

exp of a quantity *affine* in $\bar{X}$, i.e. $K \propto \exp(-\beta\bar{X})$, the same functional form the substrate side carries. The Gaussian of step 1 is thereby *excluded*: it correlates the bits and breaks the product structure (Computation 111), so the wave-function route is unobstructed in contrast to the Jacobian route. Note that the product structure fixes only the *form* (log-linear in $\bar{X}$), not the coefficient: the value $\beta_{\mathrm{KO}} = 2$ is the binary-distinction gap $2 = 1 - (-1)$, supplied separately (Computation 112), not by the factorisation.

3. *The field-strength's substrate-side value is the embedding norm of the symmetric Higgs vacuum.* For the Higgs as an emergent projected field $\phi = \Pi(\psi)$, the field-strength is its norm inside the substrate Hilbert space under the modal (tempered) inner product, $Z_\phi^2 = \mathbb{E}_\mu[\exp(-\beta_{\mathrm{KO}}\bar{X})] \to \mathrm{e}^{-1}$. Two readings of the field-strength are possible, and the selection between them is *structural*. The standard effective field theory (EFT) reading (the two-point residue of $\bar{X}$) is its susceptibility $\mathrm{Var}_\mu(\bar{X}) = 1/(4D) \to 0$. That vanishing is the *law of large numbers*: $\bar{X}$ is an intensive average of the i.i.d. bits (P1), so it self-averages to a sharp value, and a quantity the law of large numbers (LLN) sends to zero cannot be a field-strength. Reading it as one gives $\lambda_{\mathrm{SM}} \to 0$, a *free* emergent Higgs. The embedding norm is the $O(1)$ object $\to \mathrm{e}^{-1}$, the only candidate that normalises a canonical emergent field. Since the emergent EFT below M_* is the *interacting* Standard Model (Proposition F4), an interacting Higgs is required, hence $Z_\phi^2 = O(1)$ and hence the embedding-norm reading. That reading is forced by P1's LLN and F4, not by the measured value of λ_{SM}, which is thereby demoted from an input to the *test* of the prediction $b\,\mathrm{e}^{-1}$ (Computation 117). Physically, below the LG threshold $\langle\bar{X}\rangle$ is a sharp classical condensate (susceptibility $\to 0$) and the physical Higgs is the radial excitation, normalised by the embedding norm, not by the order parameter's vanishing susceptibility. The embedding-norm reading holds precisely when the substrate→Higgs reduction is *total* (every substrate mode folded into the single Higgs sector), which step 4 establishes.

4. *The reduction is total: a unique-soft-mode theorem.* Past the Landau–Ginzburg threshold (P3) the order parameter $\bar{X} = (1/D)\sum_a C_a$ is the unique soft mode, with a spectral gap (the binary gap $= M_*$) to every other substrate mode. This is Theorem F10 below; it establishes that the reduction is total, so the field-strength of step 3 is the *full* substrate embedding norm $\to \mathrm{e}^{-1}$, not a projection of it. F10 underwrites the substrate-side value; it does not by itself supply the SM-side identification.

12.5.3. What the four steps establish, and what remains

Steps 1–4 fix (B)'s type (a multiplicative field-strength renormalisation, not an additive Jacobian) and its substrate-side value ($Z_\phi^2 \to \mathrm{e}^{-1}$, total reduction proved by F10), with an $O(1/D)$ finite-size residual that vanishes in the substrate limit (the same limit in which e^{-1} itself is defined). The matching ratio $Z^2 = \lambda_{\mathrm{SM}}(M_*)/b$ has, besides this field-normalisation piece, a vertex piece $Z_{\Gamma 4}$ (the four-point 1PI correction); it is trivial to all orders, $Z_{\Gamma 4} = 1$ (Computation 119). This is the $n = 4$ instance of the all-orders S_D-equivariant decoupling lemma (Computation 118), of which F10's quadratic (Hessian) statement is the $n = 2$ instance. The quartic claim takes as *given* F10's qualitative conclusion, that $\bar{X}$ is the unique soft mode, so the non-collective modes are the gapped UV modes integrated out at the matching, together with $V = \Phi(\bar{X})$ (Theorem F8, the order parameter is the collective coordinate); what it does *not* require is F10's quadratic-order Hessian *method*. Because $V = \Phi(\bar{X})$ depends on the linear collective coordinate $\bar{X} = (1/D)\sum_a C_a$ alone, its nth derivative tensor is the all-ones tensor $\Phi^{(n)}/D^n$, which every non-collective mode v of the S_D standard representation ($\sum_a v_a = 0$) annihilates at each order n; the $n = 4$ contraction is precisely the quartic 1PI vertex, so the integrated-out substrate UV modes have zero matrix element with the Higgs quartic and generate no vertex correction (Computation 119 verifies the $n = 4$ contraction numerically to $\sim 10^{-50}$). In fact the decoupling is closed not merely as a formal all-orders lemma but *exactly and non-perturbatively*: because $\bar{X}$ is the empirical mean of the i.i.d. bits (F8 with P1), it is a *sufficient statistic* for any interaction of the form $\Phi(\bar{X})$, so the configurations at fixed

$\bar{X} = k/D$ are counted by the interaction-independent multiplicity $\binom{D}{k}$ and the substrate partition function factorises, at every finite D, as

$$Z[\Phi] \;=\; \mathbb{E}_{\mu}\Big[\mathrm{e}^{-\Phi(\bar{X})}\Big] \;=\; \sum_{k=0}^{D} \binom{D}{k} 2^{-D}\, \mathrm{e}^{-\Phi(k/D)} \tag{146}$$

(Computation 121). The transverse (non-collective) modes enter *only* through the Φ-independent multiplicity $\binom{D}{k}$, so $Z_{\Gamma 4} = 1$ (and every n-point Φ-vertex correction) vanishes *exactly*, by Fisher–Neyman factorisation [151] rather than by a Gaussian-order argument; this answers the "higher-order" concern that the quadratic F10 does not by itself reach the quartic vertex. The S_D-equivariant Morse–Bott lemma (Computation 118) is the de Moivre–Laplace shadow [101] of (146) (the all-ones derivative tensor is its leading term), and the only $D \to \infty$ limit is the *value* of the one-dimensional collective integral, whose normalisation gives e^{-1} with the standard $O(1/D)$ correction, the substrate-limit caveat every PST quantity carries, now isolated to the value rather than the structure. Consistency with a finite matched coupling is then manifest: the entire Higgs quartic at M_* is b times the field-strength dressing e^{-1}, with no substrate-UV contribution to it, so $\lambda_{\mathrm{SM}}(M_*) = b\,\mathrm{e}^{-1}$ is finite and well-defined. The matching's *internal renormalisation is therefore complete*: $Z^2 = \mathrm{e}^{-1}$, with both the field-strength (Computation 117) and the vertex (Computation 119) derived from the same decoupling. What the steps do *not* supply is a derivation of the matching on the *SM side*: the equality $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z^2$ is Wilsonian threshold matching at M_* (Propositions F4–F7), and its connection to the running SM value $\lambda_{\mathrm{SM}}(M_*) \approx 0.0930$ rests on F4, the emergence of the SM as the EFT below M_*, the framework's central result, established in §18.1 and §13.7, not a (B)-specific gap, together with the 1.2% RGE confirmation (the collective Higgs self-loops are the SM running below M_*, which confirms $b\,\mathrm{e}^{-1} = 0.0920$, not a $Z_{\Gamma 4}$ correction). (B) is therefore *reduced*: no internal step of the matching is open, and what remains is F4 plus the empirical RGE test. The matched scaling $\Lambda^2 = D$ is the one definitional identification shared with the substrate side.

12.5.4. Substrate-side closure; the open matching step

The substrate-side content is closed in the substrate limit, as established by the four-step reduction above and Theorem F11. The tensor-product factorisation (Computation 100) forces the cutoff function to $f(x) = \exp(-x)$, and at the matched scaling $\Lambda^2 = D$ binomial product convergence delivers $(1/2^D)\,\mathrm{Tr}\, e^{-\Delta/D} \to e^{-1}$. Bridge premise (B), eq. (145), is thereby reduced to a well-typed field-strength-renormalisation identity with substrate-side value $Z^2 \to \mathrm{e}^{-1}$. What remains open is the identification of this substrate field-strength with the SM-measured coupling ratio $\lambda_{\mathrm{SM}}(M_*)/b$, a Wilsonian threshold-matching step (F4–F7) resting on the emergence premise F4 and carried, evidentially, by the 1.2% agreement.

12.6. The framework in full: F1–F11

The definitions, propositions, and theorems summarised above are stated here in full, with their proofs and established-in references, before the two main theorems.

12.6.1. Definitions (F1–F3)

12.6.1.1. Definition F1 (PST partition-function spectral triple). A PST partition-function spectral triple is a tuple $(\mathcal{A}, \mathcal{H}, D_{\mathrm{sub}}, \mu, \beta)$ where $(\mathcal{A}, \mathcal{H}, D_{\mathrm{sub}})$ is the substrate spectral triple of §1.5 with $\mathcal{A}$ the Boolean $*$-algebra on D, $\mathcal{H} = \ell^2(\mathcal{P}(D), \mu)$, D_{sub} the canonical Boolean Dirac operator; $\mu = \bigotimes_{a \in D} \mathrm{Bern}(\frac{1}{2})$ is the Bernoulli substrate measure; and $\beta > 0$ is the temperature parameter.

12.6.1.2. Definition F2 (partition function of an observable). For an observable $O\colon \mathcal{P}(D) \to \mathbb{R}$, the β-partition function of O is

$$Z(O, \beta) \;:=\; \mathbb{E}_\mu\big[\exp(-\beta\, O(C))\big]. \tag{147}$$

12.6.1.3. Definition F3 (Matched-scaling exponent). For the substrate spectral triple, the matched-scaling exponent is

$$\beta \;:=\; 2, \tag{148}$$

equal to the spectral range of a single binary distinction, $2 = 1 - (-1)$: each exchange τ_i is a self-adjoint involution ($\tau_i^2 = 1$) with eigenvalues $\{+1, -1\}$, so $(1 - \tau_i)$ has nonzero eigenvalue 2, forced with no freedom by P1's binary alphabet (equivalently the one-bit Clifford $\mathrm{Cl}(1, 0)$ range; Computation 100, 102). The numerically equal $\mathrm{KO}_{\mathrm{total}} \bmod 8 = 2$ (with $\mathrm{KO}_{\mathrm{total}} = 4 + 6 = 10$, §1.5) shares the value but is a gauge- and spacetime-sector fact, never an input here; the name "binary-gap temperature" β_{KO} is retained below as a label only. This $\beta = 2$ is also the calibration point of the Bernoulli-MGF route (the value, and only the value, for which $\mathbb{E}_\mu[\exp(-\beta \bar{X})] \to e^{-1}$): e^{-1} is the signature of the binary gap.

12.6.2. Propositions (F4–F7)

These restate ingredients already established elsewhere in the book; they are collected here so the framework is self-contained.

12.6.2.1. Proposition F4 (PST EFT below M_*). *The PST effective field theory below the matched scale $M_* = 4\pi m_h \sqrt{2/3}$ equals the Standard Model exactly: no PST modes propagate below M_*, so the SM Lagrangian (gauge + Higgs + matter) is the leading effective content of the theory below M_*.*

Established as the leading-order Wilsonian EFT in §12.2, equation (127); the gauge + Higgs + matter content is derived as the projected spectral action in Chapter 13.

12.6.2.2. Proposition F5 (Wilsonian matching at M_*). *The IR SM couplings at M_* equal the UV PST couplings transformed by the standard Wilsonian matching condition, expressed as the substrate-side path-integral over modes with momenta in $[M_*, \Lambda]$ where $\Lambda = \sqrt{D}$ is the matched-scaling*

substrate cutoff.

Established in standard EFT machinery; the substrate-side specification is given in F8 below.

12.6.2.3. Proposition F6 (PST UV Higgs quartic). *The PST UV theory at M_* has bare Higgs quartic coupling $\lambda_{\mathrm{PST}}^{\mathrm{bare}} = b = 1/4$, fixed by the Landau–Ginzburg modal potential coefficient.*

Established in §1.5; cf. equation (140).

12.6.2.4. Proposition F7 (Substrate UV-IR scale separation). *At matched scaling $\Lambda = \sqrt{D}$, the substrate UV modes lie in the interval $[M_*, \Lambda]$. The path-integral over these modes contributes the substrate vacuum-fluctuation factor to the SM EFT below M_*.*

Established by matched-scaling A1 of §18.1.

12.6.3. Theorems (F8–F11)

These carry the substrate-side content that is new to this framework.

12.6.3.1. Theorem F8 (substrate Higgs Hamiltonian). *The substrate Higgs Hamiltonian $H_{\mathrm{Higgs}}\colon \mathcal{P}(D) \to \mathbb{R}$ is uniquely determined by the three structural constraints (i) additivity (Bernoulli factorisation), (ii) matched scaling (per-site energy $1/D$), and (iii) uniformity (P1 distinguishability symmetry); explicitly,*

$$H_{\mathrm{Higgs}}(C) \;=\; \frac{1}{D}\sum_{a\in D} B_a(C) \;=\; \bar{X}(C). \tag{149}$$

No free parameters.

Proof: Computation 89; cf. equation (137).

12.6.3.2. Theorem F9 (tempered Bernoulli partition function). *The partition function of H_{Higgs} at the matched-scaling binary-gap temperature β_{KO} satisfies*

$$Z(H_{\text{Higgs}}, \beta_{\text{KO}}) \;=\; \mathbb{E}_{\mu}\big[\exp\big(-\beta_{\text{KO}}\,\bar{X}\big)\big] \;\xrightarrow{D\to\infty}\; \mathrm{e}^{-1}. \tag{150}$$

Proof: standard Cramér–Bernoulli asymptotic [150]; $\mathbb{E}_{\mu}[\exp(c\bar{X})] \to \exp(c\,\mu_{\text{site}}) = \exp(c/2)$ delivers e^{-1} at $c = -\beta_{\text{KO}} = -2$ (the $\mathrm{Cl}(1,0)$ binary-distinction gap, not Bott periodicity). Computation 88; cf. equations (135), (136).

12.6.3.3. Theorem F10 (unique soft mode; total reduction). *Let the substrate action be $\Delta + V(\bar{X})$, where $\Delta = \sum_i (1 - \tau_i)$ is the ultralocal Boolean Laplacian (P1: independent bits) and V is the Landau–Ginzburg potential, a function of the linear collective coordinate $\bar{X} = (1/D)\sum_a C_a$ alone (Theorem F8: $H_{\text{Higgs}} = \bar{X}$). Then, past the LG threshold (P3), $\bar{X}$ is the unique soft mode: every non-collective substrate mode (the S_D standard representation in the $|S| = 1$ shell, and all $|S| \geq 2$ Walsh shells) retains the free binary-gap mass ≥ 2 up to $O(1/D)$, while V shifts only the S_D-singlet direction $\bar{X}$. Hence below M_* the single surviving light field is the Higgs, the substrate→Higgs reduction is total, and the field-strength of step 3 of the four-step reduction of bridge premise (B) is the full substrate embedding norm.*

Proof (all-orders S_D-equivariant Morse–Bott lemma): because $\bar{X}$ is *linear* in the bits, the chain rule gives, at every order n, $\partial^n V/\partial C_{a_1}\cdots\partial C_{a_n} = V^{(n)}(\bar{X})/D^n$, the totally symmetric all-ones tensor, independent of the indices. Under S_D the single-bit-fluctuation space $\mathbb{R}^D$ splits into the singlet line $\mathbb{R}|1\rangle$ (carrying $\bar{X}$, $|1\rangle = \sum_a e_a$) and the standard representation $W = \{v : \sum_a v_a = 0\}$ (the non-collective modes). An all-ones tensor contracted in any slot with $v \in W$ gives $(\sum_a v_a)(\cdots) = 0$, so V has zero matrix element on W at every order: the non-collective modes are free fields with mass set by Δ alone (the binary gap ≥ 2). S_D-equivariance (Schur's lemma [158]) guarantees the singlet and W never mix and that no S_D-invariant interaction can deform the W-

block. Thus the Hessian, to all orders, is Morse–Bott: a soft critical direction (the order parameter $\bar{X}$, softened to criticality at threshold) and a uniformly gapped normal bundle W. The only coupling between bits beyond this is the $O(1/D)$ sum-constraint correlation, which vanishes as $D \to \infty$, so the gap is protected non-perturbatively in the substrate limit. (Computations 115, 116; formalised as the all-orders lemma in Computation 118, and closed exactly and non-perturbatively by sufficiency of the empirical mean in Computation 121, of which this lemma is the Gaussian shadow.) □

12.6.3.4. Main theorem of the framework.

12.6.3.5. Theorem F11 ($Z^2 \to \mathrm{e}^{-1}$ in the substrate). *In the substrate $D \to \infty$ limit, with the substrate-side field-strength factor closed (steps 1–4, including Theorem F10) and the Wilsonian matching $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z_\phi^2$ of Propositions F4–F7, the PST framework delivers the asymptotic identification*

$$\lambda_{\mathrm{SM}}(M_*) \;=\; b \cdot Z(H_{\mathrm{Higgs}}, \beta_{\mathrm{KO}}) \;\xrightarrow{D\to\infty}\; \frac{\mathrm{e}^{-1}}{4} \;=\; \frac{1}{4\mathrm{e}} \;\approx\; 0.0920, \tag{151}$$

equivalent to

$$Z^2 \;=\; \frac{\lambda_{\mathrm{SM}}(M_*)}{b(M_*)} \;\xrightarrow{D\to\infty}\; \mathrm{e}^{-1}, \tag{152}$$

with the substrate-side fact $Z(H_{\mathrm{Higgs}}, \beta_{\mathrm{KO}}) = \mathrm{e}^{-1}$ structurally derived from P1 ($\mu_{\mathrm{site}} = 1/2$), the binary-distinction gap $\beta_{\mathrm{KO}} = 1 - (-1) = 2$, and the matched-scaling limit, with no free parameters.

Proof: combining propositions F4 (PST EFT = SM below M_*), F5 (Wilsonian matching), F6 ($\lambda_{\mathrm{PST}}^{\mathrm{bare}} = b$), and F7 (UV scale separation) with theorems F8 ($H_{\mathrm{Higgs}} = \bar{X}$) and F9 (binary-gap-tempered partition function $\to \mathrm{e}^{-1}$), the four-step reduction of (B) (steps 1–3 and Theorem F10) yields equation (151), conditional on the Wilsonian matching $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z_\phi^2$ identifying the substrate field-strength with the SM-side coupling ratio. □

12.6.4. Remark on what the framework achieves

Theorem F11, together with the four-step reduction of (B), derives the *substrate-side* value $Z_\phi^2 \to \mathrm{e}^{-1}$ from P1–P3 in the substrate limit, in the language of a partition-function-level analogue of the Chamseddine–Connes correspondence [49, 50] adapted to the discrete substrate triple. Every *substrate-side* input is derived from PST postulates. The framework therefore:

- reduces the matching premise (B) to a field-strength renormalisation identity. The identity is well-typed (multiplicative, not the additive Jacobian of Computation 103) and its substrate-side value e^{-1} is derived within PST rather than inherited from an SM-side CC postulate, via the wave-function-renormalisation route (steps 1–3) and the unique-soft-mode theorem (F10);
- separates the substrate-side closure (F8, F9, F10, exact in the substrate limit) from the substrate→SM-coupling matching (resting on the emergence premise F4 + the RGE test, supported numerically at 1.2%), which remains the single open step; the matching's internal renormalisation is fully derived (field-strength and vertex $Z_{\Gamma 4} = 1$, Computation 119), the all-orders decoupling closed exactly and non-perturbatively by sufficiency of the empirical mean (Computation 121, eq. (146)), with only the standard $O(1/D)$ caveat on the *value* e^{-1} remaining on the substrate side;
- provides a unified mathematical statement of the Z^2 identification in the language of partition functions on substrate observables;
- documents what is and is not established at each step.

12.7. Status of the derivations

This section lists each derived result of the chapter with its present status, from fully established theorems to the items that remain conditional. The ledger comes first; the notes that follow give the reasoning behind each entry.

Result	**Status**
$\phi_{\text{Higgs}} = \Pi(\psi)$	Identification (§15.2)
$M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV	Derived, *scalar sector only*; the full-SM extension is settled in the negative (Computation 120)
$Z^2 \approx \mathrm{e}^{-1}$, hence $\lambda_{\text{SM}}(M_*) = b\,\mathrm{e}^{-1}$	**Reduced, not closed**: the substrate side is derived from P1–P3; the SM-side matching rests on the emergence premise F4 plus the RGE test
A7 equicoercivity	Theorem given A1, A2, A4 and A3
d_0 and the projection coefficient	Fixed by the canonical normalisation, a premise rather than a postulate: $d_0 = 2.2$ am (Computation 201)

The one entry that is not a theorem is the Z^2 row, and it is the one the chapter's headline number depends on. The notes below say why in each case.

$\phi_{\text{Higgs}} = \Pi(\psi)$ **(§15.2).** The SM Higgs field and the modal order parameter are the same field at different scales, with $m_h = 125$ GeV recovered via the LG quartic.

The modal scale M_*. $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV (Computations 67, 69, 93, 97). The apparent zero-mode off-by-one of the boson–fermion cancellation (Computations 67, 69) is closed by the substrate-vs-emergent distinction (Computations 93, 97). The substrate Walsh zero mode $\chi_\emptyset$ identifies with the vacuum expectation value (VEV) direction (the constant function on M), structurally distinct from the emergent 4-component SM Higgs doublet generated by inner fluctuations on the spectral triple $M \times F$. The resolution is therefore a VEV-vs-fluctuation split, with the projected zero mode being $\langle\phi\rangle$ rather than the physical h. Including the zero mode in N_B does not double-

count $C_H = 6\lambda = 3/2$, and the scalar self-mass relation $m_h^2 = (3/2)\, M_*^2/(16\pi^2)$ recovers as a parameter-free one-loop derivation *of the scalar sector*. The extension of the boson–fermion cancellation to the full SM (top/W/Z sectors) requires three further ingredients: the configuration-to-particle map, cross-sector coupling universality at M_*, and the inner-fluctuation sign assignment of §12.2. It is settled in the negative (Computation 120). No such EFT-layer cancellation exists within PST, consistent with the large negative Veltman residual of equation (125), and M_* is irreducibly a scalar-sector result.

$Z^2 \approx \mathrm{e}^{-1}$. The substrate-side factor $Z_H(\beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}$ is derived structurally from P1–P3, and its identification with the SM-side ratio $\lambda_{\mathrm{SM}}(M_*)/b$ (bridge premise B) is reduced within PST to a field-strength-renormalisation identity whose internal renormalisation is fully derived (field-strength and vertex $Z_{\Gamma 4} = 1$), so $\lambda_{\mathrm{SM}}(M_*) = b\,\mathrm{e}^{-1} \approx 0.0920$ rests on the emergence premise F4 plus the empirical RGE test (1.2%, 1.4σ, at three-loop precision), not a free parameter. The full account, with the spectral-action route, the Bernoulli-MGF derivation, the seven-step Wilsonian matching, and the partition-function-level Chamseddine–Connes framework, is canonical at §19.7.

A7 equicoercivity. The Mosco-convergence theorem of §18.1 is stated on seven assumptions A1–A7. A7 is a theorem of §18.1.6 under A1, A2, A4 and the equidistribution condition A3, the rate hypothesis, the closure margin of emergent spacetime. The three closure subsections (a)–(c) deliver the rate hypothesis at proof-detail level: the explicit $\Pi_n^\dagger$ construction, the thin-support d_0-corrections, and the Riemannian extension on $M = \mathbb{R} \times S^3$ (full detail in §18.2). Computations 70 and 83 corroborate the rate on the i.i.d. surrogate across single and seven-family test sets. The remaining work is explicit numerical-constant computation in the Sobolev embeddings (standard finite-element analysis, Brenner–Scott [102]). A3 is a single diffeomorphism-invariant equidistribution inequality: the leading-coefficient bound $\eta_{K,R} < d_0^2/4$, not a mixing rate (Computation 129). Its fluctuation channel is closed unconditionally (McDiarmid, Computation 125). The leading-coefficient inequality itself is supplied by the maximum-entropy sublimation of P3, by typicality: the

realised configuration is a typical one under μ, its emergence cloud equidistributed at the rate A3 states. That is a selection principle internal to P3, not a fourth postulate (no P4), and it is corroborated on the i.i.d. surrogate (Computations 70, 83).

12.8. Summary

This chapter builds the quantum fields out of the substrate and goes after one number, the Higgs quartic. The gradient term of P3, evaluated at the vacuum, gave the geometry; evaluated away from it, the same term gives the spectrum of excitations, and that branch projects into quantum field theory (§12.1). The interacting vacuum is stable at one loop and reproduces the Standard Model as its Wilsonian EFT below M_*, with the Higgs mass nominally relating to the modal scale via $m_h = M_*\sqrt{6}/(8\pi)$ at the LG quartic $\lambda_{\mathrm{PST}} = \frac{1}{4}$. The numerical value $M_* = 4\pi m_h\sqrt{2/3} \approx 1285$ GeV is a parameter-free one-loop derivation *of the scalar sector*, fixed by the single measured input m_h (the full-SM extension of the cancellation is settled in the negative, Computation 120), with the substrate-vs-emergent zero-mode distinction (Computations 93, 97) placing N_B (including the substrate Walsh zero mode) and C_H (the emergent Higgs doublet) in structurally distinct layers of the projection chain. The coherence length d_0 and the projection coefficient are Landau–Ginzburg/condensate quantities whose magnitudes are not pinned by the postulates alone, d_0 being fixed by the canonical normalisation at 2.2 am (Computation 201). The binary-gap-tempered partition function closes the substrate side of the bridge premise, $Z_H(\beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}$ from P1–P3 alone (§12.5); the matching of that factor to the measured quartic rests on the emergence premise F4 and is tested rather than proved, at 1.4σ on current inputs (Computation 202). The framework F1–F11 states what the chapter assumes and what it proves, and the status ledger records which entries are theorems and which is the one step reduced rather than closed.

13. Gauge Structure, Quantisation, and the Standard Model

The known particles are sorted by three symmetries, written SU(3), SU(2) and U(1). Nobody knows why those three and not some others. This chapter argues they were never a choice: they are symmetries the substrate already had, read off after the projection.

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

The projected field theory of §12.1 establishes a Lorentzian scalar field theory carrying a topologically quantised U(1) charge, with a massive radial mode (the Higgs) and a massless Goldstone mode. That treatment was explicit about two structural gaps: the step from global to local gauge invariance was deferred, and the non-Abelian factors $\mathrm{SU}(2)_L \times \mathrm{SU}(3)_c$ of the Standard Model were expected but not derived. The present chapter closes both gaps and establishes a third result: given the Standard-Model hypercharges, gauge anomaly cancellation is not a numerical coincidence but a structural consistency check that forces $N_c = 3$, the value PST fixes independently as 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 §13.7. The substrate's parity-selected real structure delivers the chirality grading $\gamma_K = i\omega_{\mathrm{JW}}$ exactly in the Jordan–Wigner representation (Computation 19), so the chirality projector $f = (I + i\omega)/2$ comes from the substrate itself. Furey 2014 [121] then delivers A_F as the algebra generated by $\mathrm{Stab}_{G_2}(e_7)$ on the chiral half of the minimal left ideal $\mathcal{S} = (\mathrm{Cl}(0,6) \otimes \mathbb{C})\, p$, generated by f's canonical rank-1 refinement p, Furey's primitive idempotent (an enveloping algebra, not a commutant, which on a minimal left ideal is $\mathbb{C}$; §13.7). 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. The full derivation is carried out in Computation 6 (gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$

as the unimodular unitary group of A_F; inner-fluctuation count delivering 12 gauge bosons plus the Higgs scalar). Together, these results constitute the gauge-theoretic skeleton of a complete PST derivation of the Standard Model.

13.2. Path-integral quantisation of the projected theory

The Euclidean modal functional $\mathcal{F}[\psi, C]$, equipped with the Bernoulli product measure μ derived in §2.5, defines a substrate partition function:

$$Z_{\text{sub}} \;=\; \int_{\mathcal{C}} \exp\bigl(-\mathcal{F}[\psi(C), C]\bigr)\, d\mu(C). \tag{153}$$

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

$$Z_{\text{QFT}} \;=\; \int \mathcal{D}\psi\, \mathcal{D}g_{\mu\nu}\, \exp\bigl(iS[\psi, g]\bigr), \tag{154}$$

with $S = S_{\text{EH}} + S_{\text{matter}}$ the Einstein–Hilbert plus matter action of Chapter 7. This is the path integral of quantum gravity coupled to the modal field. The substrate grain $\ell_* = \hbar c/M_* \approx 1.5 \times 10^{-19}$ m provides a physical ultraviolet cutoff $\Lambda_{\text{UV}} = 1/\ell_*$: not a calculational regularisation but a genuine breakdown scale of the continuum description. Renormalisability is the statement that low-energy observables are insensitive to ℓ_*, which holds for separations $\gg \ell_*$ exactly as the Casimir analysis requires.

Expand around the vacuum $\psi = r_0 e^{i\theta_0}$, and the quadratic fluctuation action S_2 of equation (123) generates Feynman rules for two modes: a massive scalar (mass $m_h^2 = 4\varepsilon$, equation (122)) and a massless pseudoscalar (the Goldstone mode [43]). The metric fluctuations $h_{\mu\nu}$ supply graviton propagators alongside them. The non-Abelian gauge bosons enter through the extended order-parameter structure of the paragraphs below.

13.3. 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. 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, and the gauge kinetic term

$$S_{\text{Maxwell}} \;=\; -\frac{1}{4e^2}\int_M F_{\mu\nu}F^{\mu\nu}\sqrt{-g}\;d^4x, \tag{155}$$

is supplied by the inner fluctuations of the spectral action (Computation 6), not by integrating out θ.

What integrating out θ actually gives, and which boson this is. Expanding $r_0^2 D_\mu\theta\, D^\mu\theta$ produces $r_0^2 A_\mu A^\mu$, and that term is gauge invariant *precisely because* θ is there to be eaten: $(\partial\theta - A)$ is itself invariant, and in unitary gauge $\theta = 0$ nothing is left over. This is the Stueckelberg form of the Abelian-Higgs mechanism, and it gives the connection a mass $m_A^2 = 2e^2 r_0^2$ at the condensate scale. *So the boson gauged here is not the photon.* The premise of this section is $|\psi| = r_0 \neq 0$, that is, that the U(1) acting on θ is *broken*, and a gauged broken U(1) gives a massive vector, which is a superconductor rather than electromagnetism (§9.2). Once the electroweak embedding of §13.4 is in place, this is the Z (Computation 172).

The photon is the connection of the *unbroken* combination $Q = T_3 + Y/2$, massless because Q annihilates the vacuum ((186), §13.4). That is the Standard Model's own unbroken-generator argument, which PST inherits; and there Q's annihilation of the vacuum is what *fixes* $Y(H) = +1$, so masslessness is an input to that fit rather than an independent output of it. Minimal coupling of charged fields via D_μ follows from gauge covariance of the projected kinetic terms (Computation 166). The gauging is not an additional postulate: it is the consequence of Π assigning an independent vacuum phase to each spacetime point.

13.3.1. Minimal coupling is forced, and this is the algebra

"Follows from gauge covariance" is an assertion until someone does the variation, and the variation says something sharper than the phrase does. Computation 166 carries it out with explicit Dirac matrices and a completely general spinor.

The variation is one term, and it names the current. Under $\psi \mapsto e^{i\alpha(x)}\psi$ the free Dirac Lagrangian changes by

$$\delta\mathcal{L} \;=\; -\left(\bar{\psi}\gamma^\mu\psi\right)\partial_\mu\alpha, \tag{156}$$

exactly, with no other term. The object that appears is the *vector* current, so only a counterterm built on that same current can cancel it. That is already the uniqueness statement: the coupling is not chosen from a list of gauge-invariant candidates, it is read off the variation.

And the transformation law is an output. Writing the counterterm as $-e\,(\bar{\psi}\gamma^\mu\psi)A_\mu$ and asking what shift of A_μ cancels (156) returns $A_\mu \mapsto A_\mu - e^{-1}\partial_\mu\alpha$, and returns it uniquely. Neither the form of the vertex nor the gauge transformation rule is put in by hand; both fall out of requiring the phase to be local.

The one alternative fails twice. The other dimension-four candidate is the axial coupling $-e\,(\bar{\psi}\gamma^\mu\gamma^5\psi)A_\mu$. It cannot cancel (156), because that variation

carries the vector current and this term carries a different one. And it could not source a gauge field in any case: $\partial_\mu \partial_\nu F^{\mu\nu}$ vanishes identically by antisymmetry, so $\partial_\mu F^{\mu\nu} = e\, j^\nu$ admits only a conserved source, and on shell the two currents behave differently,

$$\partial_\mu\big(\bar\psi\gamma^\mu\psi\big) = 0, \qquad \partial_\mu\big(\bar\psi\gamma^\mu\gamma^5\psi\big) = 2im\,\bar\psi\gamma^5\psi, \tag{157}$$

both verified against the explicit spinor. The axial current is not conserved for $m \neq 0$, and is therefore not available.

What this fixes is the *form* of the electromagnetic vertex, and that is all it fixes. *The strength is not supplied.* An earlier version wrote $e^{-2} \propto r_0^2$, which cannot hold: $1/e^2$ is dimensionless in four dimensions while r_0^2 carries $[\text{mass}]^2$, and in D dimensions the two carry $[\text{mass}]^{D-4}$ and $[\text{mass}]^{D-2}$, so the proportionality is impossible in *every* dimension. The combination $e^2 r_0^2$ that does carry $[\text{mass}]^2$ is the Stueckelberg mass of §13.3, not an inverse coupling. So the vertex's form is forced (Computation 166) and its magnitude is an open input; the photoelectric cross-section computed from it (Computation 189) inherits that.

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

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

$$H \;=\; \begin{pmatrix} H^+ \\ H^0 \end{pmatrix} \;\in\; \mathbb{C}^2, \qquad |H|^2 \;=\; |H^+|^2 + |H^0|^2. \tag{158}$$

The scalar ψ of Chapter 7 is the neutral component H^0 in the broken phase; H is the Standard Model Higgs doublet [75, 76].

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

$$\mathcal{V}_{\mathrm{EW}} \;=\; \left\{\, H \in \mathbb{C}^2 \;:\; |H|^2 = v^2/2 \,\right\} \;\cong\; S^3, \tag{159}$$

where $v = 246$ GeV is the electroweak vev [93] ($r_0 \leftrightarrow v/\sqrt{2}$ in the doublet convention of (201); $r_0 = v$ in the single-component normalisation of Chapter 7, the two reconciled in Chapter 15). The three-sphere S^3 is diffeomorphic to the group manifold of SU(2):

$$S^3 \;\cong\; \mathrm{SU}(2), \tag{160}$$

and its isometry group is $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$. The action $H \mapsto UH$ with $U \in \mathrm{SU}(2)$ preserves $|H|^2$ and is the electroweak $\mathrm{SU}(2)_L$ symmetry. Three of the four real degrees of freedom in H are Goldstone bosons of the breaking $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{em}}$; they are absorbed by the $W^\pm$ and Z bosons, giving them mass [77, 78, 79]. The remaining degree of freedom is the physical Higgs boson of mass $m_h^2 = 4\varepsilon$, identified in equation (122).

The $\mathrm{SU}(2)_R$ isometry is broken by the directional character of the threshold crossing: $\varepsilon = |T(C)| - \tau > 0$ is an oriented statement in which tension exceeds the threshold from below, selecting a preferred chiral handedness. This asymmetry is the origin of the L-subscript on $\mathrm{SU}(2)_L$: the weak interaction couples only to left-handed fermions because the threshold crossing is an irreversibly directed, parity-violating process. Parity violation in the weak sector is therefore not an external assumption but a theorem of the asymmetric modal threshold [83, 84].

13.4.1. Two descriptions of one would-be Goldstone phase

The single Landau–Ginzburg (LG) phase mode θ on the substrate side appears in two formally different but physically identical guises, describing the same field at the same scale in different language. *The first edition listed three and the first two were the same object counted twice*, with the phase of H^0 appearing once among the eaten Goldstones and once again as a separate $\mathrm{U}(1)_{\mathrm{em}}$ gauge

redundancy; the correction is Computation 172 and the arithmetic is that a four-component doublet cannot carry five real degrees of freedom.

1. *Substrate view* (§13.3, §9.2): θ is the phase of the complexified order parameter ψ, which P2's determination of the V_7 direction does not fix, $\hat{\tau}$ being a real direction carrying no phase (Computation 171). It is the mode whose wave equation §4.12 argues from, described before it is eaten.

2. *Electroweak doublet view* (this section): with $\psi = H^0$ enriched to the Standard Model (SM) Higgs doublet $H = (H^+, H^0)^{\mathrm{T}}$, θ *is* the phase of the neutral component, and it is one of the three angular Goldstones absorbed into the longitudinal polarisations of $W^{\pm}$ and Z under $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{em}}$. The broken diagonal combination, the Z, is the one that takes it.

   It is not eliminable by $\mathrm{U}(1)_{\mathrm{em}}$*, and the reason is the same fact that makes that symmetry unbroken*: $Q = T_3 + Y/2$ annihilates the vacuum, and it annihilates H^0 as well, H^0 carrying charge zero, so a $\mathrm{U}(1)_{\mathrm{em}}$ transformation does not move θ at all and has nothing to gauge away (Computation 172). Four real components, three eaten, one physical and massive; no massless scalar survives.

A third statement, which is about the zero mode rather than a third view of the same one. A spatially uniform rotation $\psi = r_0 \mathrm{e}^{i\omega t}$ would carry a conjugate Noether charge, that of the broken neutral direction the Z absorbs (not the electromagnetic charge, which does not move θ), and the vacuum, the energy minimum, has $\omega = 0$ (§9.1) (§9.1, where the charge is derived as $J^0 = 2r_0^2\omega$). That is consistent with θ being eaten: a mode absorbed into the longitudinal Z carries no independent zero mode either. The first edition gave the charge here as $J^0 = 4\varepsilon\omega$ and cited §16.5, which derives neither; the form and the reference now follow the section that does.

These are not different fields nor different mechanisms; the first two are one Goldstone phase described in two vocabularies appropriate to two observational regimes (the infrared (IR) and the electroweak), and the third is a statement about the zero mode rather than a third view of the same one. The unifying

statement is: θ does not propagate as an independent physical scalar in emergent spacetime. The only propagating scalar dof of the LG–Higgs sector is the radial mode, the physical Higgs h of mass $m_h \simeq 125$ GeV.

13.4.2. Relation to existing electroweak approaches

PST's substrate-Higgs mechanism is this section together with the identification of §15.2 and the gauging of §13.3. How it compares against the Connes–Chamseddine–Marcolli spectral action, Technicolor and composite-Higgs models, Coleman–Weinberg radiative spontaneous symmetry breaking (SSB) [91], Little Higgs models, and supersymmetric extended Higgs sectors is treated in §17.5.

13.4.3. The doublet does not fix chirality

The derivation above gives the $\mathrm{SU}(2)_L$ doublet but not its chirality (its action on left-handed components only). Chirality is a separate structural consequence of the directed modal threshold, treated in §14.4. The same threshold-direction parity argument that drives parity violation selects which of the two su(2) factors of $\mathrm{so}(4) = \mathrm{su}(2)_L \oplus \mathrm{su}(2)_R$ on the spatial S^3 is identified with the physical $\mathrm{SU}(2)_L$.

13.5. Colour $\mathrm{SU}(3)_c$: octonionic origin and Casimir decomposition

Colour $\mathrm{SU}(3)_c$ 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 underlying Furey-2014 foundation, the action of the octonion chain algebra $\overleftarrow{\mathbb{O}}$ on one-generation matter content delivering the SU(3) structure, is verified in Computation 16. The Standard Model gauge group is the unimodular unitary group of A_F (Chapter 19), with $\mathrm{SU}(3)_c$ the unimodular unitaries of the $M_3(\mathbb{C})$ factor, $\mathrm{SU}(2)_L$ the unit quaternions of $\mathbb{H}$, and $\mathrm{U}(1)_Y$ the surviving abelian factor. Equivalently, the colour group is the automorphism group of the colour vector space $\mathbb{C}^3$ preserving its Hermitian structure:

$$\mathrm{Aut}\big(\mathbb{C}^3,\, \langle\cdot,\cdot\rangle\big) \;=\; \mathrm{U}(3) \;=\; \mathrm{U}(1) \times \mathrm{SU}(3)_c, \tag{161}$$

with the overall U(1) removed by unimodularity. The eight gluons arise as the inner fluctuations of the Dirac operator on the $M_3(\mathbb{C})$ factor (Computation 6). This section presents the colour $\mathrm{SU}(3)_c$ argument at its three depths: the algebraic statement above, the G_2-stabiliser-of-V_7 source of the $M_3(\mathbb{C})$ factor, and the rigorous Casimir decomposition that verifies the irrep content exactly.

13.5.1. Octonionic origin: the G_2 stabiliser of a V_7 direction

The colour factor descends from the octonions. The exceptional group $G_2 = \mathrm{Aut}(\mathbb{O})$, 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 $\mathrm{SU}(3) \subset G_2$ the seven directions of V_7 split as

$$\mathbf{7} \longrightarrow \mathbf{1}_{\mathbb{R}} \oplus \mathbf{3}_{\mathbb{C}} \cong \mathbb{R} \oplus \mathbb{C}^3, \tag{162}$$

where $\mathbf{3}_{\mathbb{C}}$ is the colour triplet on which $\mathrm{SU}(3)_c$ acts. Equivalently, the complexified octonions $\mathbb{C} \otimes \mathbb{O}$ generate $\mathrm{Cl}(6) \cong M_8(\mathbb{C})$, whose minimal left ideal carries one generation's colour and electric-charge content (§14.1), and whose three fermionic ladder operators $\alpha_k = \frac{1}{2}(L_{a_k} + iL_{b_k})$ for the pairs $(a_k, b_k) = (1,3), (2,6), (4,5)$ of the complex structure J below build one generation's colour triplet (the Gunaydin–Gursey/Furey construction; Computation 16). The colour multiplicity is the complex dimension of this triplet,

$$N_c = \dim_{\mathbb{C}}(\mathbf{3}_{\mathbb{C}}) = 3, \tag{163}$$

the rank of the $M_3(\mathbb{C})$ colour 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 $\mathrm{SU}(N_c)$ gauge theory with $N_c \geq 2$, inherited by PST from the standard quantum chromodynamics (QCD) argument [81, 82, 80].

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 $\mathrm{SU}(3)_c$ 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 (§14.2, Computations 46–48). The three quaternionic sub-algebras of $\mathbb{O}$ containing the complex sub-algebra $C_{\hat{\tau}} = \mathrm{span}\{1, \hat{\tau}\}$ are identified with three generations of matter, with $\hat{\tau}$ acting within each generation as the chiral preferred direction. PST therefore derives *one* three, not two: $\hat{\tau}^{\perp}$ is three-dimensional over $\mathbb{C}$, and that single fact is the source of both the colour rank $N_c = 3$ and the generation count $N_{\mathrm{gen}} = 3$. The assignment of the three octonionic sectors to generations versus colour states is the colour-versus-generation distinguishability caveat of the octonionic-Standard-Model programme (Furey 2014 [121], Open Questions section, p. 4): on $\mathbb{O}$ alone the three sectors mix under colour SU(3) because $\mathrm{Stab}_{G_2}(\hat{\tau}) \cong \mathrm{SU}(3)$ acts on $\hat{\tau}^{\perp} \cong \mathbb{C}^3$ as the fundamental.

This book previously claimed to resolve that on the Dixon side, and the claim is withdrawn. The argument was that the ambient $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ makes the three generation sectors orthogonal (Computation 48). They are orthogonal, exactly; that is not the issue. Orthogonality of subspaces is not invariance under a group, and Computation 175 measures both facts side by side: the sectors are mutually orthogonal to machine zero *and* colour leaks out of them at O(1). Worse for the enlargement, colour acts on the ambient only through the $\mathbb{O}$ factor, as $I \otimes g$, so $\|(1 - I \otimes P)(I \otimes g)(I \otimes P)\| = \|(1 - P)\, g\, P\|$: the leakage on the 64-real-dimensional ambient equals the leakage on $\mathbb{O}$ alone, to 0. The ambient supplies room and no invariance.

And within $\hat{\tau}^{\perp}$ no such splitting can exist at all. The commutant of $\mathrm{Stab}_{G_2}(\hat{\tau})$ there is exactly $\mathbb{C}$, so by Schur the representation is the irreducible fundamental **3** and has no proper invariant subspace (Computation 175). The obstruction is therefore a theorem rather than the dimension count the register used to state it: a $\mathbf{3}_{\mathrm{colour}} \otimes \mathbf{3}_{\mathrm{gen}}$ would need nine complex dimensions against three, and irreducibility additionally forbids the direct-sum splitting the Fano lines actually realise. PST inherits this question rather than closing it; the full

account is given in §14.2.

13.5.2. Rigorous Casimir decomposition

The qualitative octonionic-origin story above admits a rigorous, fully checked algebraic version. The Lie group $G_2 = \mathrm{Aut}(\mathbb{O})$ acts on the direction algebra $V_7 \cong \mathbb{R}^7$; fixing a unit direction $e_7 \in V_7$ singles out the stabiliser $\mathrm{Stab}_{G_2}(e_7) = \mathrm{SU}(3)$, 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 Cayley convention of Computation 16 (Chapter 14 uses the Fano-plane convention of Computation 46), J pairs the basis directions as (e_1, e_3), (e_2, e_6), (e_4, e_5) (Computation 16, section 16 of the script), and the $J = +i$ eigenspace is the SU(3) triplet

$$T \;=\; \mathrm{span}_{\mathbb{C}}\{u_1, u_2, u_3\}, \tag{164}$$

$$u_1 = \tfrac{1}{\sqrt{2}}(e_1 - ie_3),\; u_2 = \tfrac{1}{\sqrt{2}}(e_2 - ie_6),\; u_3 = \tfrac{1}{\sqrt{2}}(e_4 - ie_5), \tag{165}$$

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

$$\mathbb{C}^8 \;\cong\; \langle e_0\rangle \oplus \langle e_7\rangle \oplus T \oplus \bar{T} \;=\; \mathbf{1} \oplus \mathbf{1} \oplus \mathbf{3} \oplus \overline{\mathbf{3}}. \tag{166}$$

The structural claim about §13.5's qualitative identification is then a sharp Clebsch–Gordan computation [171], fully verified in Computation 17. Lifting each Gell-Mann [167] 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 $\mathbf{8} \otimes \overline{\mathbf{8}}$, and the $(\mathbf{1} + \mathbf{1} + \mathbf{3} + \overline{\mathbf{3}}) \otimes (\mathbf{1} + \mathbf{1} + \overline{\mathbf{3}} + \mathbf{3})$ Clebsch–Gordan expansion gives

$$M_8(\mathbb{C}) \;=\; 6 \cdot \mathbf{1} \oplus 5 \cdot \mathbf{3} \oplus 5 \cdot \overline{\mathbf{3}} \oplus 2 \cdot \mathbf{8} \oplus \mathbf{6} \oplus \overline{\mathbf{6}}. \tag{167}$$

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

To verify equation (167) numerically, Computation 17 constructs the SU(3) quadratic Casimir $C_2 = \sum_a (\mathrm{ad}\,\Lambda_a)^2/4$ as a 64×64 Hermitian matrix and diagonalises it. The Casimir is an irrep discriminator: it takes the values 0 on **1**, 4/3 on **3** and $\overline{\mathbf{3}}$, 3 on **8**, and 10/3 on **6** and $\overline{\mathbf{6}}$. 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 (167) is therefore not a conjecture but an exact Clebsch–Gordan identity verified at the irrep-multiplicity level.

13.5.3. What the decomposition does and does not settle

The verified decomposition of equation (167) 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 $\mathbb{C} \otimes \mathrm{Cl}(0,6) \cong M_8(\mathbb{C})$ through the natural representation (166), 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 the algebra-identity question. It does *not* by itself fix $N_{\mathrm{gen}} = 3$ generations: equation (167) gives the colour decomposition of $M_8(\mathbb{C})$ but is silent on the generation count. The generation count is fixed separately, by the V_7-directional structure of P1–P3 and the $\hat{\tau}$-selection at modal sublimation. Three quaternionic subalgebras of $\mathbb{O}$ contain $\hat{\tau}$, and they deliver three disjoint generation sectors of $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$, each carrying one generation's Standard-Model matter content (§14.2, Computations 46–48).

13.5.4. The strong sector: what is derived and what is inherited

The preceding derivation is narrower than the phrase "colour from the substrate" suggests, and the boundary is worth drawing explicitly, because the strong sector is where a reader is most likely to expect a claim the theory does not make.

What is derived. The gauge *group* and its rank: $\mathrm{SU}(3)_c$ as the stabiliser $\mathrm{Stab}_{G_2}(\hat{\tau})$ of the threshold direction in V_7, acting on $\hat{\tau}^{\perp} \cong \mathbb{C}^3$ as the fundamental, whence $N_c = 3$ by equation (163). That is a structural result and it is where the substrate does work.

What is inherited. Everything dynamical. Asymptotic freedom, confinement, the running of α_s, the hadron spectrum and the whole apparatus of non-perturbative calculation are standard quantum chromodynamics, taken over unmodified. The derivation above says as much of confinement, which follows from the non-Abelian structure of any $\mathrm{SU}(N_c)$ theory with $N_c \geq 2$ and not from P1–P3. PST supplies the group on which QCD is built; it does not re-derive QCD, and nothing in this work should be read as claiming a substrate account of confinement. Such an account would be a separate and considerably harder result.

What the comparison is, so the boundary is not read as a deficit. No framework derives the group. The Standard Model posits $\mathrm{SU}(3)_c$ and fixes $N_c = 3$ by measurement, and grand unified theories posit a larger group that contains it. A structural derivation of the group and its rank from the threshold direction is therefore more than the incumbent offers, not less, and the dynamical half is inherited by every framework alike.

What follows for lattice results. Lattice QCD is the principal non-perturbative method in this sector, and it is worth stating plainly what its results do and do not bear on here. They are computations *within* the theory PST reproduces, not tests of the derivation above. A lattice determination that revised a hadronic quantity, or that resolved a discrepancy between lattice and data-driven methods, would revise QCD phenomenology and leave the

group-theoretic result untouched: the stabiliser calculation does not depend on any hadronic number. The corollary is equally plain, and cuts the other way: agreement between lattice QCD and experiment is not evidence for PST either.

The one channel by which a strong-sector number reaches a PST prediction, and its size. The strong coupling enters this work at a single load-bearing place, the three-loop error budget for the Higgs self-coupling test (Computation 202), where the parametric uncertainty on $\lambda_{\mathrm{SM}}(M_*)$ decomposes as m_t ± 0.00061 (its pole-mass ambiguity included), m_h ± 0.00025 and α_s ± 0.00018. The strong coupling is the *smallest* of the three, by a factor of three against the top mass. A percent-level revision to a hadronic input therefore does not move the comparison that §19.7 reports. The residual there is dominated by the top and Higgs masses, and it is those, not α_s, whose tightening will sharpen or break the test.

What would be a problem. A demonstration that the strong interaction is not described by an SU(3) gauge theory, or that the colour multiplicity is not three, would remove the structure this section derives and with it the octonionic reading of the internal algebra. Neither is an open question, which is why the derivation is a consistency result rather than a prediction: PST recovers a gauge group already established, and the weight it can carry is correspondingly limited.

13.6. Consistency of the octonionic interpretation

The colour-block embedding $\mathrm{SU}(3) \subset \mathrm{SU}(4) = \mathrm{Spin}(6)$ and the octonionic embedding $\mathrm{SU}(3) \subset G_2 \subset \mathrm{Spin}(6)$ 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 section shows that the octonionic interpretation is consistent with every other major structural result of PST.

13.6.1. Embedding-independence of the structural results

Computation 12 (together with Computation 13) compares the canonical and octonionic embeddings side-by-side under each of the verified results listed in Chapter 19. 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 §18.3 hold in both cases (Computation 9); the tree-level custodial protection $\Delta T = 0$ of §15.11 holds in both cases (Computation 10); and the chiral-bimodule selection of §14.4 holds in both cases (Computation 11).

13.6.2. 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, §14.4) 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). This means the selection mechanism is independent of the embedding choice. Numerical verification on the lepton-sector finite triple gives a residue

$$\left\|[[D_F, a],\, Jb^*J^{-1}]\right\|_F \;\sim\; 2 \times 10^{-5} \qquad \text{(Yukawa } D_F\text{, both embeddings)} \tag{168}$$

versus

$$\left\|[[D_F, a],\, Jb^*J^{-1}]\right\|_F \;\sim\; 10 \quad \text{(generic } D_F\text{, both embeddings)}, \tag{169}$$

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 §14.4.

13.6.3. Structural scope under the octonionic embedding

Computation 1's substrate-level structural-scope result, sections 1–6 of the script (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 $\mathrm{Cl}(0,6)$. A side-by-side comparison of the canonical and octonionic embeddings on the load-bearing observables is recorded in Computation 13. Computation 1, section 7 of the script (the inner-fluctuation extension) holds under the octonionic embedding (Computation 14): 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 §14.3 is preserved.

Under the octonionic interpretation, four things hold. The gauge structure of the octonionic-SU(3) closure is structurally derived (§13.5). Every established structural result of PST that depends on the embedding choice is verified under both canonical and octonionic embeddings. The order-one condition selects the Yukawa form identically. And the Yukawa-contingency reading of §14.3 is preserved. The structural derivation of $N_{\mathrm{gen}} = 3$ on top of the verified octonionic embedding follows from the V_7-valued asymmetric tension and the $\hat{\tau}$-selection at modal sublimation, closed in §14.2 (Computations 46–48).

13.7. The internal algebra A_F derived from P1–P3

This section closes Conjecture 1, stated formally as §19.5: *P1–P3, with the maximal-division-algebra (Dixon) hyperspinor envelope (a consequence of P1, Computation 168) of §1.5, imply that the internal algebra of PST's foundational object is* $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$. The closure is derived from P1–P3 with no further appeal to the Chamseddine–Connes–Marcolli classification. The argument has five pieces. The rigorous decomposition is delivered in Computations 16 and 19 below.

13.7.1. Piece 1: the chirality grading

Computation 3 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 the Standard-Model internal real K-theory (KO) dimension six of Connes' classification. Computation 19, section 6 of the script closes the identification of γ_K with the Cl(0,6) chirality grading. In the Jordan–Wigner Majorana representation [177] of D bits, the six fermion Majoranas χ_a, χ'_a $(a = 1, 2, 3)$ generate $\mathrm{Cl}(0, 6)$ on the $2^D = 8$-dimensional Hilbert space, and the volume element

$$\omega_{\mathrm{JW}} \;=\; \chi_1\chi'_1\chi_2\chi'_2\chi_3\chi'_3 \tag{170}$$

satisfies $\omega_{\mathrm{JW}}^2 = -I$. Computation 19, section 6 of the script verifies

$$\begin{aligned} \gamma_K \;&=\; Z \otimes Z \otimes Z \;=\; i\,\omega_{\mathrm{JW}} \\ &\quad (\big\|\gamma_K - i\omega_{\mathrm{JW}}\big\| = 0.0 \text{ to machine precision}). \end{aligned} \tag{171}$$

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

$$f \;=\; \tfrac{1}{2}(I + \gamma_K) \;=\; \tfrac{1}{2}(I + i\omega), \qquad \bar{f} \;=\; \tfrac{1}{2}(I - i\omega), \tag{172}$$

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

13.7.2. Piece 2: f as the canonical idempotent

Spin(6) acts on $\mathrm{Cl}(0, 6) \otimes \mathbb{C} \cong M_8(\mathbb{C})$ via the spinor representation, and on $\mathrm{Cl}(0, 6)$ by conjugation through its embedding in $\mathrm{Cl}(0, 6)^{\mathrm{even}}$ as the bivector Lie algebra. Computation 18 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 $\mathrm{Cl}(0,6)$ by grade

$$\mathrm{Cl}(0,6) \;\cong\; \mathbf{1} \oplus \mathbf{6} \oplus \mathbf{15} \oplus \mathbf{20} \oplus \mathbf{15} \oplus \mathbf{6} \oplus \mathbf{1}, \tag{173}$$

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, section 4 of the script 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 $H_- \cong \overline{\mathbf{4}}$ of $\mathrm{SU}(4) \cong \mathrm{Spin}(6)$. Hence $f = (I + i\omega)/2$ and $\bar{f}$ are the unique non-trivial Spin(6)-invariant idempotents in $\mathrm{Cl}(0,6) \otimes \mathbb{C}$: the two chirality projectors, whose canonical rank-1 refinements are Furey's primitive idempotents.

13.7.3. Piece 3: $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ from the chain algebra

Given the chirality projector f, Furey 2014 [121] works in the chain algebra $\mathbb{C} \otimes \overleftarrow{\mathbb{O}} \cong \mathrm{Cl}(0,6) \cong M_8(\mathbb{C})$ and identifies:

- generators of SU(3) within the chain algebra acting on $\mathbb{C} \otimes \mathbb{O}$, under which the eight-dim spinor space carries one generation's worth of fermion content under $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ (one SU(3) triplet, one anti-triplet, and two singlets, per equation (176); the six-triplet, six-singlet count is the decomposition of the full chain-algebra *module* $\mathcal{S}_{\text{full}}$, three minimal left ideals carrying the three generations; the decomposition of the chain algebra itself is equation (167));
- a U(1) hypercharge generator μ giving the hypercharges $-1/2$, $+1/6$, $-1/6$, $+1/2$ on the four chiral subspaces $S_{\pm\pm}$.

What Furey 2014 does *not* derive in her paper is the $\mathrm{SU}(2)_L$ factor of A_F; her "Open Questions" section (p. 4) flags this explicitly and points to the Dixon algebra $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ as the appropriate setting for the full electroweak structure.

13.7.4. Piece 4: SU(2) from the Lorentzian Dirac sector

The substrate's Mosco limit (Computation 4) delivers a 4-d Lorentzian spacetime $M = \mathbb{R} \times S^3$ with signature $(+,-,-,-)$ (one timelike + three spacelike directions; Computation 2), whose local Lorentz Clifford algebra is, by Cartan classification, $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$.[1] Dirac spinors are accordingly $\mathbb{H}^2$-valued, and right-multiplication by $\mathbb{H}$ on the spinor module is an *internal* symmetry (it commutes with the external Cl(1,3) Lorentz action); this is the classical structural observation of Dixon 2010 [146] ("isospin SU(2) arises from the subset of unit-quaternion elements of the right-$\mathbb{H}$ action"). Computation 66 verifies the four structural facts explicitly: (a) the Cl(1,3) anticommutators hold on $\mathbb{H}^2 \cong \mathbb{R}^8$; (b) right-$\mathbb{H}$ satisfies the quaternion algebra; (c) right-$\mathbb{H}$ commutes with Cl(1,3); (d) right-$\mathbb{H}$ is chirality-preserving (commutes with the Lorentz pseudoscalar). The exponentiated unit-quaternion group $\mathrm{Sp}(1) \cong \mathrm{SU}(2)$ is the internal symmetry of the emergent Lorentzian Dirac sector.

Vectorial vs chiral. The internal SU(2) delivered by Piece 4 is vectorial: by construction it commutes with the Lorentz pseudoscalar γ^5 and therefore acts identically on left- and right-handed chiral subspaces. The Standard-Model $\mathrm{SU}(2)_L$ acts chirally (on left-handed fermions only); identifying Piece 4's vectorial SU(2) with $\mathrm{SU}(2)_L$ requires a parity-breaking selection that singles out one chirality. This selection is supplied by PST's directed modal threshold (§14.4): the directed crossing $\varepsilon = |T(C)| - \tau > 0$ orients the S^3 volume form and breaks $\mathrm{SO}(4) \cong \mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$ down to the left factor, making the

[1] The signature is written $(+,-,-,-)$ to make explicit the assignment of the timelike component's sign: this is the standard Clifford-algebra convention under which $\gamma_0^2 = +I$ and $\gamma_i^2 = -I$, giving the quaternionic algebra $M_2(\mathbb{H})$. Computation 2 derives the physical content (1 timelike + 3 spacelike directions) independent of sign convention; we adopt the convention that makes the spinor sector's quaternionic structure manifest. Under the alternative convention $(-,+,+,+)$ the Clifford algebra would be $\mathrm{Cl}(3,1) \cong M_4(\mathbb{R})$, which is real and has no internal SU(2); the two conventions describe the same physics but the former is the one in which the emergent SU(2) that Piece 4 derives below is visible at the real-algebra level. The internal SU(2) is itself a physical symmetry of the emergent theory, independent of this representation choice: its existence follows from the Lorentzian signature derived in Computation 2 together with the right-$\mathbb{H}$ action of Computation 66, and the signature convention fixes only whether it is manifest at the real-algebra level, not whether it is present.

vectorial Piece 4 SU(2) chiral on the threshold-selected sector. The chiral identification therefore rests on Piece 4 (the algebraic existence of the internal SU(2)) plus the directed-threshold parity argument (the chirality selection); both are structural to PST.

13.7.5. Piece 5: Dixon hyperspinor synthesis

Dixon 2010 [146] (page 9) shows that the full Dixon algebra $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ is the spinor space of $\mathrm{Cl}(0,9) \cong M_{16}(\mathbb{C})$ and synthesises the $\mathrm{Cl}(0,6)$ chain-algebra structure of Piece 3 with the right-$\mathbb{H}$ Dirac-spinor structure of Piece 4 into a single algebra whose internal A_F is $\mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ (with the standard $\mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ identification on one generation, replicated three times via Furey's six-triplet decomposition). The Dixon algebra is the tensor of the four real normed division algebras $(\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O})$ classified by Hurwitz. Using the full tensor $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ is a *maximal-structure selection*, the same instruction as the $V_7 = \mathrm{Im}(\mathbb{O})$ choice flagged in §1.5 rather than a separate one: it is a consequence of P1 rather than a selection principle beside it, a proper sub-tensor requiring a datum P1 does not supply (Computation 168) (other division-algebra Standard-Model programmes use proper subalgebras, e.g. $\mathbb{C} \otimes \mathbb{O}$ *à la* Furey, or $\mathbb{H} \otimes \mathbb{O}$). Under this selection the synthesis lands on $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$.

13.7.6. Conjecture 1, closed (within the Dixon synthesis)

Combining Pieces 1–5:

$$\begin{aligned} &\text{P1–P3} + \text{Mosco convergence} + \text{Dixon synthesis} \\ &\quad\Longrightarrow\; A_F \;=\; \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C}) \end{aligned} \tag{174}$$

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.

3. The canonical anticommutation relations (CAR) algebra of the substrate at D bits is $\mathrm{Cl}(0, 2D)$; for $D = 3$ this is the eight-dimensional Spin(6) spinor representation, on which $M_8(\mathbb{C})$ acts by left multiplication (Computation 16).

4. In the Jordan–Wigner representation, $\gamma_K = i\omega_{\mathrm{JW}}$ exactly (Computation 19, section 6 of the script, equation (171)). Hence $f = (I + \gamma_K)/2$ is the substrate's chirality projector (of rank 4 on $H_F = \mathbb{C}^8$); Furey's primitive idempotent is its canonical rank-1 refinement, the vacuum idempotent inside the chiral half.

5. Computation 18 and Computation 19, section 4 of the script certify that f is the canonical Spin(6)-invariant choice (uniqueness up to chiral swap).

6. Furey 2014 [121] delivers $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation of fermion content via the $\mathrm{Cl}(0, 6)$ chain-algebra structure (*not* the full A_F).

7. The emergent Lorentzian Dirac sector $\mathrm{Cl}(1, 3) \cong M_2(\mathbb{H})$ delivers the internal SU(2) via right-$\mathbb{H}$ on the spinor module (Computation 66).

8. Dixon 2010 [146]'s hyperspinor framework synthesises the chain-algebra structure of (6) with the Dirac-spinor structure of (7) and identifies $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ as the internal algebra acting on the combined hyperspinor.

Step 4 (equation (171)) is the bridge from the substrate's parity grading to Furey's idempotent. Step 7 (Computation 66) is the bridge from emergent Lorentzian spacetime to the internal SU(2). The derivation invokes Dixon's hyperspinor framework rather than the Chamseddine–Connes–Marcolli axiomatic classification of finite spectral triples; the two routes arrive at the same algebra by different structural paths.

The closure rests on three pieces of standard mathematical content, all well established:

- The CAR algebra of D fermion bits is isomorphic to $\mathrm{Cl}(0, 2D)$ via the Jordan–Wigner Majorana representation.

- Furey 2014 [121]'s $\mathrm{SU}(3) \times \mathrm{U}(1)$ action on the $\mathrm{Cl}(0,6)$ chain algebra $\mathbb{C} \otimes \mathbb{O}$, partitioning the chain algebra into six SU(3) triplets and six singlets identifiable as three generations of fermion content under $\mathrm{SU}(3)_c \times \mathrm{U}(1)$.
- Dixon 2010 [146]'s hyperspinor framework $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$, in which the Cl(0,6) chain algebra and the emergent Lorentzian Dirac sector's right-$\mathbb{H}$ Sp(1) (Piece 4) combine to give the full internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$.

None is a PST-specific assumption. With PST read as a derivation of physics from P1–P3 plus the Connes-spectral-triple framework, including the emergent Lorentzian Dirac sector that supplies the $\mathbb{H}$ factor, the algebra identity is unconditional given the Dixon / maximal-tensor selection of §1.5 (the selection principle flagged above). The two structural items that accompany the closed Conjecture 1 have different status. The Dirac-convergence step is closed at the relaxed-criterion level (§7.1, Chapter 19). The three-generation count $N_{\mathrm{gen}} = 3$ is derived from the directional content of V_7, the three quaternionic sub-algebras through $\hat{\tau}$ (§14.2), and placed in Furey 2014's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0,6)$, the module where the octonionic programme reads it; the colour-versus-generation distinguishability caveat that programme carries is settled inside PST in the negative, $\hat{\tau}^{\perp}$ being the irreducible **3** of $\mathrm{SU}(3)_c$ (§13.5, §14.2, Computation 175).

The remainder of this subsection lays the closure out explicitly: the minimal left ideal of $\mathrm{Cl}(0,6)$, the commutant calculation, the emergent $\mathrm{Cl}(1,3)$ contribution, the Dixon synthesis, and the three-generation reading, step by step. The mathematical content is well-established in [121, 146, 49]; the substrate-side identifications are PST-specific.

13.7.7. $\mathrm{Cl}(0,6)$ representation theory

The real Clifford algebra $\mathrm{Cl}(0,6)$ has dimension $2^6 = 64$ over $\mathbb{R}$. Complexified, $\mathrm{Cl}(0,6) \otimes_{\mathbb{R}} \mathbb{C} \cong M_8(\mathbb{C})$, the full 8×8 matrix algebra over $\mathbb{C}$. In the standard Cayley basis $\{e_1, \ldots, e_6\}$ of $\mathrm{Im}(\mathbb{O})$, the Clifford generators $\gamma_i = L_{e_i}$ (left octonionic multiplication by e_i) satisfy $\{\gamma_i, \gamma_j\} = -2\delta_{ij}$ on the spinor module

$\mathbb{C} \otimes_{\mathbb{R}} \mathbb{O} \cong \mathbb{C}^8$ ([121], Computations 19, 16). The chirality grading

$$\omega \;=\; \gamma_1\gamma_2\gamma_3\gamma_4\gamma_5\gamma_6 \tag{175}$$

satisfies $\omega^2 = -1$, so $i\omega$ has eigenvalues ± 1 and defines the chirality projector $f = \frac{1}{2}(I + i\omega)$, of rank 4 on the spinor module $\mathbb{C}^8 = H_+ \oplus H_-$. The substrate's parity-selected real structure on an odd number of distinction bits delivers the KO-6 Connes sign triple $(+,+,-)$ exactly (Computation 3), and in the Jordan–Wigner representation $\gamma_K = i\omega$ to machine precision (Computation 19): f is the substrate's own parity-grading projector seen through $\mathrm{Cl}(0,6) \otimes \mathbb{C}$.

13.7.8. The minimal left ideal

Within the chiral half selected by f, the canonical primitive idempotent p (Furey's vacuum idempotent, the Jordan–Wigner ladder product; rank 1, with $p = fp = pf$, unique up to the inner action certified in Piece 2) generates a minimal left ideal of $\mathrm{Cl}(0,6) \otimes \mathbb{C} \cong M_8(\mathbb{C})$: $\mathcal{S} = \big(\mathrm{Cl}(0,6) \otimes \mathbb{C}\big)p$. As a $\mathbb{C}$-vector space, $\mathcal{S} \cong \mathbb{C}^8$; viewed inside $M_8(\mathbb{C})$ via the rank-1 projector p, $\mathcal{S}$ is a single column of $M_8(\mathbb{C})$, its generator lying in the chiral half selected by f.

Decomposing $\mathcal{S}$ under $\mathrm{Stab}_{G_2}(e_7) = \mathrm{SU}(3) \subset G_2 = \mathrm{Aut}(\mathbb{O})$ gives the matter-content split

$$\mathcal{S} \;=\; \mathbb{C} \;\oplus\; \mathbb{C}^3 \;\oplus\; \overline{\mathbb{C}^3} \;\oplus\; \mathbb{C}. \tag{176}$$

The two singlets are the neutrino and electron; the triplet and anti-triplet are the up- and down-quark colour triplets of one generation, with electric-charge content matching the Standard Model hypercharges (Furey 2014 §2; Computation 16).

13.7.9. Furey's chain-algebra action and the commutant

Furey's construction acts the chain algebra $\mathcal{C}$, the algebra generated by left multiplication on $\mathbb{C} \otimes \mathbb{O}$, which is $\mathrm{Cl}(0,6) \otimes \mathbb{C}$, on the minimal left ideal $\mathcal{S}$ via the chain product: $x \in \mathcal{C}$ acts as $\psi \mapsto x \cdot \psi$ in $M_8(\mathbb{C})$. The chain product is associative on $\mathcal{C}$ (it is just matrix multiplication on $M_8(\mathbb{C})$), even though octonionic multiplication on $\mathbb{O}$ alone is not.

The commutant of this action is *small*, and it is worth being exact about it. $\mathcal{S}$ is a minimal left ideal of $\mathcal{C} \cong M_8(\mathbb{C})$, hence its defining irreducible representation, so by Schur

$$\mathrm{Comm}(\mathcal{C}, \mathcal{S}) \;=\; \mathrm{End}_{\mathcal{C}}(\mathcal{C}p) \;\cong\; p\,\mathcal{C}\,p \;=\; \mathbb{C}, \tag{177}$$

one complex dimension, acting as the overall U(1) phase on $\mathcal{S}$. Nothing non-trivial commutes with the whole of $M_8(\mathbb{C})$. *An earlier version of this passage gave the right-hand side as* $\mathbb{C} \otimes M_3(\mathbb{C})$*, nine dimensions; that is withdrawn.*

$M_3(\mathbb{C})$ arises instead as an *enveloping* algebra rather than a commutant, and the distinction costs nothing because the object is the same one the chapter already uses. With $\omega = L_{e_1} \cdots L_{e_6} = L_{e_7}$ and the chirality $\gamma = i\omega$ of (175), the octonionic $\mathrm{SU}(3) = \mathrm{Stab}_{G_2}(e_7)$, the centraliser of $J = L_{e_7}$ *within* G_2, acts on the chiral half $\mathcal{S}_+ \cong \mathbb{C}^4$ of (176) as $\mathbf{1} \oplus \mathbf{3}$. The unital algebra it generates there is

$$\langle \not{\Lambda},\ \Lambda_a \rangle \;\cong\; \mathbb{C} \oplus M_3(\mathbb{C}), \tag{178}$$

ten complex dimensions, with two-dimensional commutant as Schur requires (Computations 16, 17). That is the $\mathrm{U}(1) \oplus M_3(\mathbb{C})$ part of A_F, with $\mathbb{C}$ the overall phase and $M_3(\mathbb{C})$ acting on the colour triplets $\mathbb{C}^3 \oplus \overline{\mathbb{C}^3}$. Restricting to unimodular unitaries gives $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation, with hypercharge assignments forced by the decomposition of (176) (Furey 2014 §3–4). This is the substrate-side delivery of $\mathrm{SU}(3)_c \times \mathrm{U}(1)$; $\mathrm{SU}(2)_L$ is not yet present and requires the emergent-spacetime contribution below.

13.7.10. $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ and the internal SU(2)

The Mosco limit of the substrate's Dirichlet form yields emergent spacetime $M = \mathbb{R} \times S^3$ of KO-dimension four (Computation 4; C1 of §7.1, a theorem under A1–A7 rather than under P1–P3 alone). The local Clifford algebra of M is

$$\mathrm{Cl}(1,3) \;\cong\; M_2(\mathbb{H}), \tag{179}$$

by the Cartan classification: signature $(1,3)$ gives $M_2(\mathbb{H})$ (not $M_4(\mathbb{R})$ as for signature $(3,1)$). The spinor module is $\mathbb{H}^2 \cong \mathbb{R}^8$ on which $\mathrm{Cl}(1,3)$ acts by left $\mathbb{H}$-multiplication and matrix action.

Right multiplication by $\mathbb{H}$ commutes with $\mathrm{Cl}(1,3)$ acting from the left (associativity of $\mathbb{H}$ on itself, distinct from non-associativity of $\mathbb{O}$). The right-$\mathbb{H}$ action defines an *internal* symmetry of the spinor module not part of the Lorentz group:

$$R_q \,:\, \psi \mapsto \psi \cdot q, \qquad q \in \mathbb{H}, \quad [R_q, \mathrm{Cl}(1,3)] = 0. \tag{180}$$

Restricting to unit quaternions $|q| = 1$ gives $\mathrm{Sp}(1) \cong \mathrm{SU}(2)$. Computation 66 verifies the Clifford anticommutators, centraliser, $\mathfrak{su}(2)$ generators $\{R_{iH}/2, R_{jH}/2, R_{kH}/2\}$, and chirality preservation. This is the substrate-derived SU(2).

Identifying this spacetime-side internal SU(2) with the SM electroweak $\mathrm{SU}(2)_L$ factor of A_F requires combining with the substrate-side $\mathrm{Cl}(0,6)$ chain algebra above through the Dixon hyperspinor synthesis. Chirality of the resulting $\mathrm{SU}(2)_L$ is sourced from the directedness of the modal threshold crossing, $\varepsilon = |T(C)| - \tau > 0$ (§13.4), which selects the left chiral isometry factor of $S^3 \cong \mathrm{SU}(2)$ over the right one.

13.7.11. Dixon synthesis: $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$

Dixon's hyperspinor framework [146] constructs a single 64-complex-dimensional module

$$T \;=\; \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O} \tag{181}$$

on which all three Hurwitz-classified normed division algebras act in their respective representations:

- $\mathbb{C}$ acts as scalar multiplication (overall phase, the U(1) of the Furey commutant).
- $\mathbb{H}$ acts by right-multiplication, generating the internal $\mathrm{Sp}(1) \cong \mathrm{SU}(2)$ from

the emergent Dirac sector.

- $\mathbb{O}$ acts by left-multiplication through the Furey chain algebra, generating $\mathrm{SU}(3)_c \times \mathrm{U}(1)$.

The associativity needed for each action to define a representation follows from each algebra acting on its own tensor factor: the non-associativity of $\mathbb{O}$ is sequestered to the $\mathbb{O}$ factor, while the $\mathbb{H}$ and $\mathbb{C}$ actions commute with it. Adopting the full Dixon envelope $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ rather than $\mathbb{O}$ alone is the *maximal-tensor selection principle* of §1.5.

13.7.12. The algebra identity $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$

The full internal algebra of the Connes Standard-Model spectral triple is the unimodular-unitary group of the algebra A_F acting on the Dixon hyperspinor T via the actions above. Tracking each contribution:

- The $\mathbb{C}$-action contributes the $\mathbb{C}$ summand of A_F (overall U(1)).
- The right-$\mathbb{H}$ action contributes the $\mathbb{H}$ summand of A_F (internal SU(2) on the Dirac sector).
- The left-$\mathbb{O}$ chain-algebra action contributes the $M_3(\mathbb{C})$ summand of A_F (the enveloping algebra gives $M_3(\mathbb{C})$ by (178)).

Together,

$$A_F \;=\; \mathbb{C} \;\oplus\; \mathbb{H} \;\oplus\; M_3(\mathbb{C}). \tag{182}$$

The unimodular unitary group of A_F is $\mathrm{U}(1) \times \mathrm{SU}(2) \times \mathrm{SU}(3)$, the Standard-Model gauge group. Up to the SU(2)-doublet hypercharge assignment (forced by photon masslessness, §13.4) and the chirality of $\mathrm{SU}(2)_L$ (sourced by the directedness of the threshold crossing, $\varepsilon > 0$, §13.4), this is the full SM gauge content on one generation.

13.7.13. Three generations from six SU(3) triplets

Furey 2014 §5 observes that the full chain-algebra module decomposes under SU(3) into six colour triplets, one generation being one triplet and one anti-triplet, and reads the six as three generations. In PST the count is not

read from the module but derived from the directional content of V_7: the three quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$ index three disjoint per-generation sectors, each carrying the sixteen chiral states of one Standard-Model generation (Computations 46–48), and colour and generation are one three rather than two, $\hat{\tau}^{\perp}$ being the irreducible **3** of $\mathrm{SU}(3)_c$ (Computation 175). The argument, its inputs and its scope (what is structural, what the Yukawa and Cabibbo–Kobayashi–Maskawa (CKM) structure leaves to $T(C)$, and why six of the eight gluons are generation-changing, Computation 106) are given once in §14.2 and §14.3.

13.7.14. The gauge group as substrate automorphisms

Combining the colour factor with the $\mathrm{U}(1)_Y$ and $\mathrm{SU}(2)_L$ factors derived from the vacuum topology (§13.4), the unimodular unitary group of A_F is the Standard Model gauge group:

$$G_{\mathrm{SM}} \;=\; \mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y, \tag{183}$$

as established by Glashow, Weinberg, and Salam [79, 77, 78].

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] (§8.6).
- **Electromagnetism**, the $\mathrm{U}(1)_{\mathrm{em}}$ gauge theory of the *unbroken* combination $Q = T_3 + Y/2$ left by the electroweak breaking on $\mathcal{V}_{\mathrm{EW}} \cong S^3$ (§13.4); the photon is its connection on the U(1) bundle over M. Gauging the Goldstone phase of $\mathcal{V} \cong S^1$ by itself gives a massive vector, the Z (§13.3).
- **The weak force**, the $\mathrm{SU}(2)_L$ gauge theory from the $S^3 \cong \mathrm{SU}(2)$ vacuum manifold; left-chiral because the threshold crossing is directional; $W^{\pm}$ and Z masses from the Higgs mechanism on $\mathcal{V}_{\mathrm{EW}}$.
- **The strong force**, the $\mathrm{SU}(3)_c$ gauge theory from the $M_3(\mathbb{C})$ colour factor of the internal algebra A_F; eight massless gluons; confinement from the

non-Abelian running coupling.

13.8. Gauge anomaly cancellation as a structural check

A gauge theory is quantum-mechanically consistent only if its triangle anomalies vanish [85, 86]: 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 $\mathrm{U}(1)_Y$ 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. The structural derivation of these values is inherited from the Chamseddine–Connes–Marcolli framework [50] applied to the substrate-derived spectral triple. Given $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ (Computation 6), two ingredients fix the action of the $\mathbb{C}$-factor's U(1) on each chirality block: the chiral bimodule support delivered by the directed modal threshold (Computations 7, 8, 11), and the order-one condition. The unimodularity constraint of Computation 6 then removes the overall trace. The remaining freedom is a single integer (the unimodularity assignment selecting which $\mathrm{U}(1) \subset \mathrm{U}(A_F)$ is gauged), and the SM hypercharge pattern is the unique choice consistent with anomaly cancellation under $N_c = 3$. The structural derivation walked through in Chamseddine–Connes–Marcolli (CCM) 2007 inherits into PST through PST's structural delivery of A_F, the chiral bimodule, and the order-one condition. All six hypercharges are determined explicitly below in §13.9, from Gell-Mann–Nishijima, photon masslessness, colour confinement, Yukawa gauge-invariance, and anomaly cancellation. There the Gell-Mann–Nishijima relation and the empirical nucleon charges enter as inputs rather than PST outputs. 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, the value PST fixes independently.

Express the Standard Model fermion content of one generation as left-handed Weyl spinors [162], 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:

$$\begin{aligned}\big[\mathrm{U}(1)_Y\big]^3 = 0: \quad & N_c\Big[2\big(\tfrac{1}{3}\big)^3 + \big(-\tfrac{4}{3}\big)^3 + \big(\tfrac{2}{3}\big)^3\Big] + 2(-1)^3 + (2)^3 \\ & = \; -2N_c + 6, \end{aligned} \tag{184}$$

$$\big[\mathrm{SU}(2)_L\big]^2\big[\mathrm{U}(1)_Y\big] = 0: \quad N_c \cdot \tfrac{1}{3} + 1 \cdot (-1) \;=\; \tfrac{N_c}{3} - 1. \tag{185}$$

Both vanish if and only if $N_c = 3$, so the following holds as a theorem (the colour factor is §13.5's):

> *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 reproduced by PST structure in §13.9, given the Gell-Mann–Nishijima relation and the empirical nucleon charges as inputs.*

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

13.9. $\mathrm{U}(1)_Y$ hypercharges compatible with PST structure

Four structural facts fix all six Standard Model hypercharges with no free hypercharge parameter beyond the $\mathrm{U}(1)_Y$ normalisation. (The Gell-Mann–Nishijima relation and the empirical nucleon charges enter as inputs to the determination, as the closing paragraph makes explicit.)

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

$$Q \;=\; T_3 + \tfrac{Y}{2}, \tag{186}$$

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

(ii) Higgs hypercharge from photon masslessness. After the $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{em}}$ 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 (186) gives

$$Y(H) \;=\; 2\big[Q(H^0) - T_3(H^0)\big] \;=\; 2\big[0 + \tfrac{1}{2}\big] \;=\; +1. \tag{187}$$

(iii) Quark charges from $\mathrm{SU}(3)_c$ *confinement.* Colour confinement (§13.5) 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

$$2Q_u + Q_d = 1, \quad Q_u + 2Q_d = 0 \;\Longrightarrow\; Q_u = \tfrac{2}{3}, \quad Q_d = -\tfrac{1}{3}. \tag{188}$$

Applied to the left-handed quark doublet $Q_L = (u_L, d_L)$ via equation (186):

$$Y(Q_L) \;=\; 2\big[Q(u_L) - T_3(u_L)\big] \;=\; 2\big[\tfrac{2}{3} - \tfrac{1}{2}\big] \;=\; +\tfrac{1}{3}. \tag{189}$$

(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 $\mathrm{U}(1)_Y$ singlets (total charge $Y = 0$):

$$-Y(Q_L) + Y(H) + Y(d_R) = 0 \qquad \Rightarrow\; Y(d_R) = -\tfrac{2}{3}, \tag{190}$$

$$-Y(Q_L) - Y(H) + Y(u_R) = 0 \qquad \Rightarrow\; Y(u_R) = +\tfrac{4}{3}, \tag{191}$$

$$-Y(L_L) + Y(H) + Y(e_R) = 0 \qquad \Rightarrow\; Y(e_R) = Y(L_L) - 1. \tag{192}$$

The remaining unknown $Y(L_L)$ is fixed by the $[\mathrm{SU}(2)_L]^2[\mathrm{U}(1)_Y]$ anomaly condition (equation (185)), which requires $N_c Y(Q_L) + Y(L_L) = 0$. With $N_c = 3$:

$$Y(L_L) \;=\; -N_c\, Y(Q_L) \;=\; -3 \cdot \tfrac{1}{3} \;=\; -1, \qquad Y(e_R) \;=\; -2. \tag{193}$$

The complete Standard Model hypercharge assignment is therefore

$$\begin{aligned} &Y(Q_L) = +\tfrac{1}{3}, \quad Y(u_R) = +\tfrac{4}{3}, \quad Y(d_R) = -\tfrac{2}{3}, \\ &Y(L_L) = -1, \quad Y(e_R) = -2, \quad Y(H) = +1. \end{aligned} \tag{194}$$

Substituting into equation (184) gives $-2N_c + 6 = 0$ at $N_c = 3$, a non-trivial cross-check that no value was inserted by hand; Computation 136 evaluates all six anomaly coefficients as functions of N_c and confirms anomaly freedom forces $N_c = 3$.

Given PST's representation content and $N_c = 3$, the six SM hypercharges follow from the standard determination: the Gell-Mann–Nishijima relation $Q = T_3 + Y/2$, photon masslessness (fixing $Y(H) = +1$), the observed proton and neutron charges (fixing the $\mathrm{U}(1)_Y$ normalisation via $Q_u = +\frac{2}{3}$, $Q_d = -\frac{1}{3}$), and gauge-anomaly cancellation together with Yukawa gauge-invariance

(fixing the remainder). The non-trivial PST content is that its representation content is anomaly-consistent and reproduces the SM hypercharge pattern with no free hypercharge parameter beyond this normalisation; the Gell-Mann–Nishijima relation and the empirical nucleon charges are inputs to the determination, not PST outputs. Standard Model hypercharge quantisation is thus shown compatible with, and reproduced by, PST structure, rather than postulated independently. This closes the gauge group derivation for the bosonic sector. The one-generation fermion-sector representation content is determined explicitly in §14.1, matching the hypercharge assignments derived above; the generation count $N_{\text{gen}} = 3$ is read structurally in §14.2; and the Yukawa hierarchy and Cabibbo–Kobayashi–Maskawa (CKM) matrix are addressed by the structural-scope theorem of §14.3.

13.10. The Standard Model Lagrangian as leading-order projection

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

$$\mathcal{L}_{\text{SM}} \;=\; -\tfrac{1}{4} F^a_{\mu\nu} F^{a\,\mu\nu} + \bar{\psi}\, i \not{D}\, \psi + |D_\mu H|^2 - \mathcal{F}(H) + \mathcal{L}_{\text{Yukawa}}, \tag{195}$$

where $F^a_{\mu\nu}$ runs over all gauge factors of G_{SM}, $\mathcal{F}(H)$ is the modal potential (77) 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.

13.11. Summary

This chapter 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 §13.7 (Computations 3, 19, 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 the Standard Model hypercharge assignments: the same $N_c = 3$ that the algebra delivers also makes the quantum theory self-consistent. The generation count $N_{\text{gen}} = 3$ is read structurally from Furey 2014's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0, 6)$, with the consistent V_7-directional reading via the three quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$. The ambient chain-algebra space $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ decomposes into a generation-blind core plus three disjoint per-generation sectors, so the sectors that overlap on $\mathbb{O}$ alone become orthogonal there (§14.2, Computations 46–48).

In terms of what is derived versus chosen: the internal-algebra results are derived theorems, with $\mathrm{Cl}(0, 6) \to \mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation (Furey 2014 chain algebra, Computation 19), $\mathrm{Cl}(1, 3) \cong M_2(\mathbb{H})$ delivering the internal SU(2) via right-$\mathbb{H}$ multiplication (Computation 66), and $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ as the Dixon-synthesis combination. The Dixon envelope $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ is not a modelling choice: Computation 168 derives the extremal selection from P1 (§1.5). The colour-versus-generation distinguishability caveat of the octonionic-Standard-Model programme is not resolved by the Dixon synthesis: inside PST it is settled in the negative, $\hat{\tau}^{\perp}$ being the irreducible **3** of $\mathrm{SU}(3)_c$ (§14.2, Computations 106, 175). The Higgs quartic is not this chapter's; its account

is §19.7.

14. Fermion Content: Generations, Yukawa and Chirality

Every matter particle comes in three versions, identical in every way except mass, and nobody knows why three. This chapter argues the number is counted by the substrate's own directional structure rather than being a brute fact to be accepted, and then takes on the harder question of where the wildly different masses come from.

14.1. Fermionic representation content under Π

The Jordan–Wigner identification of §11.2 and the $G_2 \to \mathrm{SU}(3) \times \mathrm{U}(1)$ decomposition of §13.5 together fix the projection Π on the fermionic Hilbert space $\mathcal{H}_D^-$ 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 per-Q_k $\mathrm{Cl}(0,2)$ chiral structure underlying this decomposition is verified in Computation 47: each of the three quaternionic sub-algebras Q_k through $\hat{\tau}$ delivers a $4+4$ chiral split with $\omega_k^2 = -I$, structurally identical across k. 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 $\mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, decomposes as

$$\bigwedge^1 \mathbb{C}^7 \Big|_{\Pi}^{1\text{ gen}} = \begin{array}{l} (\mathbf{3},\mathbf{2})^L_{1/3} \;\oplus\; (\mathbf{3},\mathbf{1})^R_{4/3} \;\oplus\; (\mathbf{3},\mathbf{1})^R_{-2/3} \\ \oplus\, (\mathbf{1},\mathbf{2})^L_{-1} \;\oplus\; (\mathbf{1},\mathbf{1})^R_{-2} \;\oplus\; (\mathbf{1},\mathbf{1})^R_{0}. \end{array} \tag{196}$$

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 (196) reproduce the assignments derived independently in equation (194) 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. Nature exhibits three generations. Taking three copies of equation (196) gives the complete Standard Model matter content plus three right-handed neutrinos:

$$\mathcal{H}_D^-\Big|_\Pi = \bigoplus_{k=1}^{3} \Big[Q_L \oplus u_R \oplus d_R \oplus L_L \oplus e_R \oplus \nu_R \Big]_k . \tag{197}$$

Total chiral state count: $3 \times (15 + 1) = 48$ Weyl fermions (the standard $45 = 15 \times 3$ Standard Model [SM] chiral states plus three right-handed neutrinos). The three-generation count is placed in Furey 2014 [121]'s six-SU(3)-triplet decomposition of $\mathrm{Cl}(0, 6)$, whose six triplets, six singlets and their antiparticles are identified as three generations of quarks + leptons in the standard octonionic-SM reading. The $\mathrm{SU}(2)_L$ doublet pairing is supplied by the Dixon synthesis (§19.5). A separate combinatorial structure on V_7 is exhibited alongside it: the three quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$ (§14.2). It is consistent with Furey's identification, subject to the colour-versus-generation distinguishability caveat, which PST inherits rather than resolves (§13.5, Computation 175).

What this does and does not predict. Equation (197) predicts the matter content and the hypercharges, both structural. It does *not* predict the Yukawa hierarchy or the Cabibbo–Kobayashi–Maskawa (CKM) matrix. At leading order the three generations in equation (197) are interchanged by the residual SU(3)-flavour symmetry of the G_2 structure (the symmetry that permutes the three complex directions of $\mathbf{3}_\mathbb{C}$). 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.

The underlying reason (the $S_3 \subset S_D \times \mathbb{Z}_2$ invariance of the Bernoulli measure, the enumerated sub-leading mechanisms it rules out, and the resulting leading-order diagonal-equal-eigenvalue Yukawa and identity CKM) is established as a no-go theorem in §14.3; the observed values must arise from the precausal configuration C itself.

SM-sector δm_h^2 cancellation: structural and coupling-weighted layers. The structural equal-weighting $N_B = N_F$ of $\mathcal{H}_D^+$ and $\mathcal{H}_D^-$ under the Bernoulli measure (§11.2) is a closed result on the substrate-side spectral action: μ assigns the same total weight to the bosonic and fermionic sub-Hilbert-spaces, so the spectral-action-level quadratic divergence in m_h^2 from the substrate DoF count vanishes structurally.

PST's structural claim and the numerical SM Veltman bracket [94] $\mathcal{V}(M_*) := 6\lambda + \frac{9}{4}g_2^2 + \frac{3}{4}g_1^2 - 6y_t^2$ (the standard one-loop quadratic-sensitivity combination of equation (124), here evaluated with couplings run to M_*) are therefore not equivalent statements: the former is a substrate-side DoF-counting identity at the spectral-action level, the latter is an emergent-SM coupling-weighted one-loop combination at the matched scale $M_* \approx 1285\,\mathrm{GeV}$. Empirically $\mathcal{V}(M_*) \approx -3.3$ at one loop. Equation (125) evaluates the same quantity to ≈ -3.5 at electroweak-scale couplings; the difference is the running of the couplings between v and M_* under the renormalisation-group equations [RGE]. That numerical value reflects the SM-side renormalised couplings at M_* (run from observed infrared [IR] values via SM RGE), not a violation of the structural identity. The substrate-side structural cancellation sits at the same layer as the partition-function-level Chamseddine–Connes correspondence of §12.5; the SM-side Veltman bracket numerical value is a downstream coupling-weighted quantity in the projected effective field theory (EFT), governed by SM RGE truncation in the same way that $\lambda_{\mathrm{SM}}(M_*) \approx 0.0930$ is. The two layers are related by the projection Π but not identified: the structural DoF-counting identity remains a closed substrate-side theorem, and the SM-side coupling-weighted bracket is a separate phenomenology evaluation that

depends on the specific SM renormalised couplings at M_*.

14.2. 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 $N_{\text{gen}} = 3$ is read structurally from two mutually consistent constructions: (a) Furey 2014 [121]'s six-SU(3)-triplet decomposition of $\mathrm{Cl}(0,6)$ (the standard octonionic-SM identification of three generations of one-generation content), and (b) the three quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$ developed in this section. Both readings rely on the same octonionic structure; their structural compatibility with a colour-singlet generation label (i.e. the statement that a colour SU(3) rotation does not mix the three generations) is a known open caveat of the octonionic SM programme (Furey's own Open Questions [121], p. 4). The Dixon synthesis (§19.5, Computation 48) does not resolve it: in the ambient $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ the three generation sectors are orthogonal as index sets, but orthogonality is not invariance, and a colour rotation carries one sector into another (Computation 106, Computation 175). On $\mathbb{O}$ alone they overlap on a generation-blind core (Computation 47), which reflects $\mathrm{SU}(3)_{\mathrm{c}} = \mathrm{Stab}_{G_2}(\hat{\tau})$ acting irreducibly on $\hat{\tau}^{\perp} \cong \mathbb{C}^3$.

14.2.1. 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 $G_2/\mathrm{SU}(3) \cong S^6$ 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 $C_{\hat{\tau}} = \mathrm{span}\{1, \hat{\tau}\}$. For $\hat{\tau} = e_1$, these are

$$Q_1 = \{1, e_1, e_2, e_3\}, \quad Q_2 = \{1, e_1, e_4, e_5\}, \quad Q_3 = \{1, e_1, e_6, e_7\}, \tag{198}$$

   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 $\mathrm{Cl}(0,2)$ 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.

14.2.2. 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 (§13.7), decomposes under the $\hat{\tau}$-selection as

$$\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O} \;=\; \underbrace{\mathbb{C} \otimes \mathbb{H} \otimes \mathrm{span}\{1, \hat{\tau}\}}_{\textbf{core},\ \dim=16} \;\oplus\; \bigoplus_{k=1}^{3} \underbrace{\mathbb{C} \otimes \mathbb{H} \otimes \mathrm{span}\{e_{a_k}, e_{b_k}\}}_{G_k,\ \dim=16}. \tag{199}$$

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 (§14.1). The chiral structure is identical across the three $\mathcal{H}_k$, recovering the observed pattern of three generations with identical quantum numbers.

14.2.3. Colour and generation: two guises of one three

The number 3 appears in PST in two guises, which this subsection shows to be one:

- *Colour* $N_c = 3$: the rank of the $M_3(\mathbb{C})$ factor of A_F (§13.5). It arises *within* a single generation, from the $\mathrm{SU}(3) \subset G_2$ 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.

What the shared origin does and does not buy. That the two descend from one selection event makes them rigid rather than independent. Both are read off the single six-dimensional space $\hat{\tau}^{\perp}$: colour as the SU(3) acting on it as $\mathbb{C}^3$, the generation count as the three Fano lines through $\hat{\tau}$ that partition it. The two cannot be varied separately, and a theory that fixed each from its own source would have no reason to expect them to agree.

It is worth being exact about how much that is worth, because it invites an overclaim. It buys no observational signature, for two reasons this section has already supplied. First, the construction is G_2-equivariant: the count of 3 does not depend on which unit vector of S^6 is selected, and a different $\hat{\tau}'$ yields a G_2-rotated copy of the same Φ with the rotation an unphysical relabelling (§4.10). Nothing observable distinguishes one pinned direction from another, which is the same argument that empties the residual Z_2 choice of content. Second, the rigidity carries no free parameter: N_c and N_{gen} are fixed by the Hurwitz and G_2 algebra before any physics enters, so there is no slack in which a correlation could be measured. What the shared origin delivers is a structural identity, not a prediction. Nor does it extend to a third structure: the chirality of $\mathrm{SU}(2)_L$ is oriented by the scalar $\varepsilon > 0$ and not by $\hat{\tau}$ (§14.4), so it does not join this list.

The structural reading of $N_{\text{gen}} = 3$ via equation (199) is verified numerically in Computations 46–48, consistent with Furey 2014's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0,6)$ (the standard octonionic Standard Model (SM) identification). The algebraic foundation underlying the chain-algebra space $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ is established in Computations 43 and 44 (Dixon-algebra foundation: the octonion chain algebra on $\mathbb{R}^8$ and its $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ extension carrying one full generation of SM matter content). The colour-vs-generation distinguishability of the three sectors on $\mathbb{O}$ alone is non-trivial: $\mathrm{Stab}_{G_2}(\hat{\tau}) \cong \mathrm{SU}(3)$ acts on $\hat{\tau}^{\perp} \cong \mathbb{C}^3$ as the fundamental, so the three sub-algebras mix under colour SU(3), and *that mixing is not undone by the ambient algebra.* Passing to $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ makes the three sectors orthogonal as index sets (Computation 48), but orthogonal is not invariant: tensoring a colour-rotated line with factors colour does not touch leaves it colour-rotated. Computation 106 measures the leakage directly and finds it $O(1)$, with only the maximal torus $\mathrm{U}(1)^2 \subset \mathrm{SU}(3)_c$ preserving all three lines. There is a dimension obstruction too: independent colour and generation factors would need $\mathbf{3} \otimes \mathbf{3}$ inside the octonionic slot, real dimension at least 18, while $\hat{\tau}^{\perp}$ has real dimension 6. *Colour and generation are not two threes here; they are the same three*, which is the octonionic

programme's own open question rather than something PST closes. Inside this theory the question is settled in the negative: $\hat{\tau}^{\perp}$ is the irreducible **3**, so no colour-invariant generation splitting exists in it at all (Computation 175, Appendix C.4); what PST inherits is where the programme puts generations if not there, and that is not carried as a structural item of this book. Two refinements would sharpen the reading: (i) demonstrate that each G_k carries the explicit $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ representation content matching the one-generation decomposition of equation (196); (ii) establish consistency of the $\hat{\tau}$-determination $\hat{\tau} = T(D)/|T(D)|$ with the Mosco-limit projection $\Phi : S \to G$. Both are concrete operator-algebra tasks of the same character as the section-round-trip and Dirac-aware-lift closures of §18.3.

14.3. Yukawa and CKM: the structural-scope theorem

The two-layer distinction. PST separates its outputs into two kinds. The *structural* layer is whatever follows from the postulates P1–P3 alone, and is therefore the same in every instantiated reality: the gauge group, the three-generation count, chirality, spin-statistics. The *contingent* layer is whatever depends on which particular precausal configuration C actually sublimated, encoded in its asymmetric-tension distribution $T(C)$. The contingent layer plays the role that initial conditions play elsewhere in physics: real content of the world, but not derivable from the laws (§20.7 gives the quantitative comparison with the Standard Model's free-parameter count). The observed Yukawa hierarchy (charged-lepton masses in the ratio $1 : 206 : 3477$ at M_*) and the non-trivial Cabibbo–Kobayashi–Maskawa (CKM) mixing matrix belong to the contingent layer. That placement is proved, not assumed: this section states the theorem, established in Computation 1, and closes the escape routes.

The no-go theorem that forces the placement (substrate level). The Bernoulli product measure $\mu = \bigotimes_a \mathrm{Bern}(1/2)$ on $\mathcal{P}(D)$ is invariant under the full permutation group $S_D \times \mathbb{Z}_2$ acting on the D substrate bits, and generation permutation $S_3 \subset S_D$ (permuting the bit-triples assigned to the three generations) is a subgroup of it. Any observable computed from the structural

ingredients (D, δ, μ) alone must therefore be S_3-symmetric across generations. Computation 1 verifies this numerically for the leading proposed generation-asymmetric mechanisms: the Yukawa overlap integral, the multi-instanton Atiyah–Singer index [178] per representation, and the Jordan–Wigner expectation in the even-parity sector are all S_3-invariant. The consequence: at leading structural order the three Yukawa eigenvalues per sector are equal and the CKM matrix is the identity, so the observed hierarchy and mixing cannot come from the structural layer. The only remaining source is $T(C)$.

Why "by construction". This is not a failed derivation reclassified as scope: the symmetry of the structural ingredients makes generation asymmetry unreachable from them in principle, so the mass spectrum was never in scope for P1–P3. Two closures secure the placement. First, at the substrate level, the proposals for generating the hierarchy from structure are examined and shown not to break S_3 invariance. Three are considered: the V_7-direction constraint (Computation 76), the dynamical per-Q_k directed modal threshold (Computation 77), and structural-β candidates in the Boltzmann [183] mechanism. The last of these runs through the structural candidates of Computation 78, sharpened in Computation 79 to the closed form $\beta = 28 + 1/4 = 28.25$ above $\dim \mathrm{SO}(8) = 28$, with the Spin-related p_1 closure in Computation 80. Each candidate inherits an S_3-invariant Bernoulli-measure trace and cannot lift the leading-order generation degeneracy on its own. Second, the spectral-triple extension: the theorem survives the Connes–Chamseddine framework via inner-fluctuation analysis. Since A_F acts on H_F as $a^{(1)} \otimes I_{\mathrm{gen}}$ 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, a direct corollary of the bimodule structure (Computation 1, section 7 of the script).

Two supporting checks. Quantitatively, the observed flavour parameters number about twenty for Dirac neutrinos: nine Yukawa eigenvalues, four CKM parameters, and seven Pontecorvo–Maki–Nakagawa–Sakata [PMNS]. They

consume about 31% of the $2^D = 64$ real degrees of freedom available in $T(C)$ at $D = 6$ substrate bits (Computation 1, section 8 of the script; a six-bit toy, two bits per generation, where the matter-sector statements of Chapter 13 are at $D = 3$). 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. Empirically, Computation 138 checks the measured CKM matrix against the form the theorem requires, identity plus a small $T(C)$-contingent perturbation: the deviation is Cabibbo-sized, every off-diagonal entry is small, and the magnitudes are unitary to sub-percent. A genuinely order-one off-diagonal CKM matrix would have refuted the placement; the observed matrix conforms to it.

The honest scope. 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.

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

The chirality of the weak gauge group (that $\mathrm{SU}(2)_L$ acts on left-handed doublets only) is a structural consequence of PST's substrate combined with the directed modal threshold of Chapter 4. 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 $\mathrm{so}(4) = \mathrm{su}(2)_L \oplus \mathrm{su}(2)_R$, equivalently the algebra of quaternion left and right multiplications on $\mathbb{H}$. This furnishes *two* candidate 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). 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 su(2) embeddings is realised in nature?

The answer is the parity-violation mechanism PST already commits to in Chapter 4: 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.

Which quantity does the orienting. The orientation just used is supplied by $\varepsilon = |T(C)| - \tau > 0$, which is a *scalar*: an oriented inequality recording that tension exceeds the threshold from below. It is not supplied by the substrate direction $\hat{\tau} = T(D)/|T(D)|$ of P2, a V_7 vector fixed by D alone, which is the quantity that fixes colour and the generation count (§14.2). The two are different objects and the distinction is not cosmetic: whether the vector and the scalar orientation are the same datum is exactly the consistency of the $\hat{\tau}$-determination with the Mosco-limit projection Φ, which §14.2 does not establish. Until it is established, chirality is not licensed as a third structure descending from $\hat{\tau}$, and §14.2 draws the consequence.

Handedness convention. The threshold orientation predicts the existence of a chirality, not the conventional "left" vs. "right" label: $L \leftrightarrow R$ is a relabelling under spatial parity, and the assignment that the chiral factor coincides with the SM's $\mathrm{SU}(2)_L$ is a convention fixed by matching to the observed weak-interaction handedness. The mechanism predicts chirality; the empirical sign of the threshold direction picks the handedness label.

14.5. Summary

This chapter settles what the fermion sector carries structurally and what it leaves to the configuration. The representation content of one generation under Π follows from the canonical anticommutation relation (CAR) identification and the internal algebra (§14.1). The three-generation count is read from

the directional content of V_7: $\hat{\tau}^{\perp}$ decomposes into three quaternionic subalgebras through $\hat{\tau}$, and colour and generation are one three, not two (§14.2, Computation 175). The Yukawa couplings and the CKM matrix lie in the configuration's content by the structural-scope theorem (§14.3): substrate-level S_3 invariance forces S_3-symmetric Yukawa eigenvalues at leading order, so the observed hierarchy is $T(C)$-contingent rather than structural. The chirality of $\mathrm{SU}(2)_L$ is a structural consequence, the empirical sign of the threshold direction picking the handedness label (§14.4).

15. Grounding d_0 and the Casimir Prediction

Everything up to here either followed from the starting rules, matched physics already known, or put one number against a measurement already made. Here the theory sticks its neck out. It names a length, says exactly what that length would do to a force you can measure in a laboratory, and then works out how big the effect is: far too small for any instrument now built. That is still a prediction. If a measurement ever finds the effect with the wrong shape, or finds it large, the theory is wrong. A theory that cannot be wrong is not worth much.

This chapter is where the theory is put at risk. Everything before it either follows from P1–P3, recovers physics already known to be true, or, in Chapter 12, puts one number against a measurement already made; what follows here is a number that could turn out to be wrong, and an experiment not yet made that could show it. The scale d_0 is grounded first, as the Landau–Ginzburg coherence length of the modal condensate, seventeen orders above the Planck length and fixed by the canonical normalisation at 2.2 am (Computation 201), and then tied to the Standard-Model sector through the electroweak identification of the modal energy M_*. The d^{-6} Casimir signature follows, with its parameter-free coefficient $\xi = 90/\pi^2$, its kernel-shape sensitivity, and the diagnostic bound $d_0 \lesssim 2.3$ nm that the data allow. The chapter closes on the experimental reach: what a measurement at each plate separation would be sensitive to, the falsifiability programme, the Arkani-Hamed–Dimopoulos–Dvali (ADD) and electroweak-precision comparison, and the correction budget.

15.1. Modal sublimation and the substrate scales

The functor $\Phi : \mathbf{S} \to \mathbf{G}$ maps the precausal substrate to an instantiated geometry with four large dimensions (three spatial, one temporal), established by the coinstantiation argument of Chapter 5 and the dimension count $n = 4$ of §1.5. The remaining substrate structure is carried by the finite noncommutative internal factor (Chapter 19), 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,

$$d_0 \;=\; \sqrt{\frac{c_{LG}}{2\,|a_{LG}|}}, \tag{200}$$

which sets the substrate discreteness scale that enters the Casimir correction. The second is the modal scale $\ell_* = \hbar c/M_*$, the length associated with the energy M_* at which the modal condensate forms (identified below with $M_* = 4\pi m_h\sqrt{2/3}$). 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; its numerical value is not determined by the three postulates alone but by the normalisation of the gradient term, and under the canonical normalisation of §12.2 it is $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am (Computation 201). (That is the single-neutral-component normalisation used throughout Chapter 7. The doublet map of equation (201) instead reads $r_0 \leftrightarrow v/\sqrt{2}$, the same $\sqrt{2}$ that distinguishes the neutral-component vacuum from the full-doublet vev, exactly parallel to the quartic $\lambda_{\rm PST} = \frac{1}{4}$ versus $b = \frac{1}{2}$.)

15.2. The 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 central identification: the modal condensate and the Higgs condensate are the same field, viewed at different scales. The PST order parameter ψ and the Standard Model Higgs field ϕ are not two structurally analogous objects; they are one object viewed from two ends of a single energy hierarchy.

Substrate-level viewpoint (looking outward from the precausal layer). The Landau–Ginzburg dynamics of ψ on the substrate's configuration space $\mathcal{C}$ exhibit a sombrero potential $F[\psi, \varepsilon] = -\varepsilon|\psi|^2 + \frac{1}{2}|\psi|^4$ (equation (77)), with vacuum manifold $\mathcal{V} \cong S^1$ at $|\psi| = r_0$, a massive radial mode, and a massless Goldstone phase θ.

Spacetime-level viewpoint (looking inward from emergent geometry). The Standard Model Higgs field ϕ has potential $V(\phi) = -\mu^2|\phi|^2 + \lambda|\phi|^4$, vacuum manifold at $|\phi| = v/\sqrt{2}$, a physical Higgs h of mass m_h, and three angular Goldstones eaten by $W^\pm$, Z under the $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_{\mathrm{em}}$ breaking (§13.4).

These two viewpoints describe *the same sombrero*: the modal-sublimation projection $\phi = \Pi(\psi)$ carries the substrate's ψ into emergent spacetime as the Standard Model (SM) Higgs field, with

$$\varepsilon \;\leftrightarrow\; \mu^2, \qquad b \;\leftrightarrow\; \lambda, \qquad r_0 \;\leftrightarrow\; v/\sqrt{2}. \tag{201}$$

The present subsection writes the sombrero for the single complex neutral component, $F[\psi, \varepsilon] = -\varepsilon|\psi|^2 + \frac{1}{2}|\psi|^4$, a normalisation in which the quartic coefficient of $|\psi|^4$ is $\frac{1}{2}$. The doublet-convention quartic of §1.7, the coefficient of $(\Phi^\dagger\Phi)^2$, $\lambda_{\mathrm{PST}} = \frac{1}{4}$, is the same physical coupling expressed for the full doublet bilinear. The map $b \leftrightarrow \lambda$ holds within whichever convention is fixed; the field-normalisation ratio $Z^2 = \lambda_{\mathrm{SM}}(M_*)/\lambda_{\mathrm{PST}}$ of §12.2 carries the substrate-to-SM kinetic-normalisation factor and is convention-independent. The substrate-level Goldstone θ is the $\mathrm{U}(1)_{\mathrm{em}}$ phase of the neutral Higgs component H^0 in the broken phase (with $\psi = H^0$; §13.4, §13.3). The identification is fixed by the numerical match below, not assumed.

In short: PST does not introduce a Higgs field on top of the substrate. The Higgs field is what the modal condensate looks like once one is inside instantiated geometry; the substrate-level ψ is what the Higgs looks like once one views it from the precausal end. The two formalisms differ in starting language but describe one physical object across the energy hierarchy

$\Lambda_{\rm IR} \ll v \ll M_* \ll M_{\rm Planck}$.

The identification places three distinct energy scales near the condensate, and their consistency is explicit. The radial-mode mass derived in Chapter 12 (equation (122)), $m_h^2 = 4\varepsilon$, together with the vacuum radius $r_0^2 = \varepsilon/(2b)$ of equation (78), gives $m_h^2 = 8b\,r_0^2$. The three quantities are therefore not independent: fixing any two fixes the third through the quartic coefficient b. Identify r_0 with the Higgs vacuum expectation value $v = 246$ GeV, and use the measured Standard-Model Higgs self-coupling $\lambda_{\rm SM} \approx 0.13$. The convention-independent observable relation $m_h = \sqrt{2\lambda_{\rm SM}}\,v \approx 125$ GeV then returns the measured Higgs mass [87, 88], with no additional input. (The substrate bare quartic in the doublet convention is $\lambda_{\rm PST} = \frac{1}{4}$. The ratio $Z^2 = \lambda_{\rm SM}(M_*)/\lambda_{\rm PST}$ is the field-normalisation factor of §12.2, not the infrared [IR] observable $\lambda_{\rm SM}$ itself.) The modal scale $M_* = 4\pi m_h\sqrt{2/3} \approx 1285$ GeV is a fourth and separate quantity (§12.2): 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.

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

The choice of M_* is settled: it is derived as a structural prediction of PST (§12.2) from the scalar-sector self-mass relation $m_h^2 = (3/2)M_*^2/(16\pi^2)$, with $m_h = 125.25$ GeV as the sole measured input; the matching of the Landau–Ginzburg (LG) quartic to the SM Higgs self-coupling at that scale is then its test, not its source (§19.7) (the 2022 value; the 2026 Review's 125.13 GeV moves M_* by 0.1 percent).

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. At that M_*:

$$\ell_* \;=\; \frac{\hbar c}{M_*} \;=\; \frac{197.3\ \text{MeV}\cdot\text{fm}}{1285\ \text{GeV}} \;\approx\; 1.535\times10^{-19}\ \text{m} \tag{202}$$

15.2.1. Relation to other approaches to the Higgs mechanism

How the $\phi = \Pi(\psi)$ identification situates against existing literature on electroweak symmetry breaking, Connes–Chamseddine–Marcolli spectral action, Technicolor and composite Higgs, Coleman–Weinberg radiative spontaneous symmetry breaking (SSB), Little Higgs and pseudo-Goldstone Higgs, supersymmetric extended Higgs sectors, is treated in §17.5.

15.3. Grounding the scale d_0: the coherence length

Before computing the coefficient ξ and bounding the amplitude, the scale d_0 itself must be placed. The amplitude of the correction (208) is proportional to d_0^2, so its detectability turns entirely on the magnitude of d_0. Chapter 8 related G and d_0 via $G = \alpha_G d_0^2$ in natural units (equation (72)), with $\alpha_G = 6\pi/(c\, r_0^2)$ carrying the stiffness c, so the relation alone does not place d_0. A reader might assume $\alpha_G \sim 1$, which would place $d_0 \approx \ell_P \approx 1.6\times10^{-35}$ 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 does not follow from the postulates, and treating ℓ_P as the only relevant length scale in the theory is unwarranted. The scale d_0 is not the microscopic grain of the substrate but the coherence length of the modal condensate, and what fixes it is the normalisation of the gradient term: under the canonical normalisation on which every number of Chapter 12 rests, $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am, seventeen orders above ℓ_P and a factor 14.5 above the grain (Computation 201). Earlier printings placed d_0 in the nanometre range by a near-threshold argument that §15.3.1 now corrects, and read the Casimir correction as a detectable signal; at the derived value it is exact in form and 10^{-19} of the force at 20 nm, and what the chapter predicts is stated in §15.5.

15.3.1. Two distinct scales: grain versus coherence length

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

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

$$\xi_{\mathrm{BCS}} \;=\; \frac{\hbar v_F}{\pi\,\Delta} \tag{203}$$

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

PST's modal condensate has the same architecture. The substrate carries a microscopic scale ℓ_*, the intrinsic grain of the distinction space D at which primitive differences are registered. Nothing in the PST axioms forces $\ell_* = \ell_P$; and even if it did, d_0 would not equal ℓ_*. The scale d_0 is defined by the projection operator Π as the finest resolution at which the realisation map $\rho : D \to M$ can distinguish spacetime points. This is the spatial scale over which the modal order parameter ψ varies: its coherence length. In the Landau–Ginzburg functional of Chapter 4:

$$F[\psi] \;=\; \int \mathrm{d}^4x \left[\tfrac{1}{2}\, c_{LG}\, |\nabla\psi|^2 \;+\; a_{LG}\, |\psi|^2 \;+\; b_{LG}\, |\psi|^4\right] \tag{204}$$

the coherence length of ψ is [8]:

$$d_0 \;=\; \xi_\psi \;=\; \sqrt{\frac{c_{LG}}{2\,|a_{LG}|}} \tag{205}$$

Near the modal threshold, $a_{LG} \propto (T/\tau - 1)$, so $\xi_\psi \to \infty$ as $T \to \tau^+$: the correlation length diverges at the transition, exactly as in any second-order phase transition. The universe exists because $T > \tau$. But how far above τ the substrate sits is not free: $|a_{LG}| = m_h^2/4$ is pinned by the Higgs mass (§15.2), so at a given gradient normalisation the coherence length is fixed, $d_0 m_h = \sqrt{2c_{LG}}$, and the near-threshold reading earlier printings gave here, which put d_0 in the nanometre range, was silently letting $|a_{LG}|$ float to 10^{-18} of its value, a radial mode of 10^{-7} GeV in place of the Higgs (Computation 201). The superconducting analogy above quantifies the same way: $\xi/a \sim E_F/\pi\Delta$ runs to three to six orders because Δ/E_F is exponentially small, while here the analogue of the gap is m_h and the analogue of the bandwidth is M_*, whose ratio this book derives, $M_*/m_h = 4\pi\sqrt{2/3} = 10.26$ (130). So under canonical normalisation $d_0/\ell_* = \sqrt{2}\,M_*/m_h = 14.5$: the two lengths are distinct by one order of magnitude, not by ten. The Standard-Model identification of ℓ_* with the electroweak modal scale M_*, and the resulting collider-sector consistency, are developed in §15.10.

15.4. Vacuum phenomena and the Casimir effect

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

$$P_{\text{standard}}(d) = -\frac{\pi^2 \hbar c}{240\, d^4} \tag{206}$$

At $d = 100$ nm this evaluates to approximately -13 Pa, measurable experimentally [12, 13]. The effect has been measured to an experimental error rising from 0.19% at 160 nm to 9.0% at 750 nm, at 95% confidence [212, 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 realisation map $\rho : D \to M$ becomes sensitive to boundary geometry. At separations $d \gg d_0$ the two descriptions are equivalent; at $d \sim d_0$ they diverge.

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

$$P_{\mathrm{PST}}(d) = P_{\mathrm{standard}}(d) \left[1 + \xi \left(\frac{d_0}{d} \right)^2 + \mathcal{O}\left(\left(\frac{d_0}{d} \right)^4 \right) \right] \tag{207}$$

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

$$\delta P_{\mathrm{PST}}(d) = -\frac{\xi\, \pi^2 \hbar c\, d_0^2}{240\, d^6} \tag{208}$$

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. The falsifiability is two-sided but asymmetric: the d^{-6} *exponent* is a parameter-free prediction, while the *amplitude* ξd_0^2 carries the coherence scale d_0, fixed by the normalisation premise at 2.2 am (Computation 201) and bounded by the data at 2.3 nm, so a null result at any reachable precision is what the theory expects and tightens only the data's bound. A confirmed d^{-6} term is signature-distinguishable from the retarded Casimir–Polder d^{-7} force, though it does not by itself select PST from other minimal-length or effective field theory (EFT) cutoff models. This two-layer structure is developed in full

below. 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.

15.5. Casimir detectability at each separation regime

What is evaluated below, and what it is. The grid below is evaluated at the diagnostic upper bound $d_0 = 2.3$ nm, which is what the data allow. The theory itself, under the canonical normalisation of §12.2, puts d_0 nine orders of magnitude lower, at 2.2 am (Computation 201), where the correction is 1×10^{-19} of the force at 20 nm and 2×10^{-21} at 160 nm. So the grid is the reach of the *bound*, not the size of the predicted signal: it says what a measurement at each separation would be sensitive to, and a null result at any of these precisions is what the theory expects. A detected d^{-6} component at nanometre amplitude would not confirm this book; it would refute the canonical normalisation on which M_* and the Higgs match rest, since it would need $c_{LG} \sim 10^{18}$.

Evaluating at the diagnostic upper bound $d_0 \approx 2.3$ nm (the theory fixes the d^{-6} form and the coefficient $\xi = 90/\pi^2 \approx 9.12$ exactly; $d_0 \lesssim 2.3$ nm is the bound from current nanometre-Casimir data under the substrate's Gaussian projection kernel, Computation 73), the Casimir correction of equation (208) gives a concrete signal at each accessible plate separation. Figure 4 plots the prediction across the whole accessible range; the two separations worked out below are the points marked on it.

At $d = 160$ nm, the shortest separation of the existing precision measurements [212]:

$$\left(\frac{d_0}{d}\right)^2\Bigg|_{d=160\,\mathrm{nm}} = \left(\frac{2.3}{160}\right)^2 \approx 2.07 \times 10^{-4} \tag{209}$$

The PST correction with $\xi = 90/\pi^2$ is $9.12 \times 0.0208\% \approx 0.19\%$. This sits right at the measured precision floor at that separation, 0.19% at 95% confidence [212]; the theory is consistent with existing data, and a coherence length larger than the bound would already have shown up there.

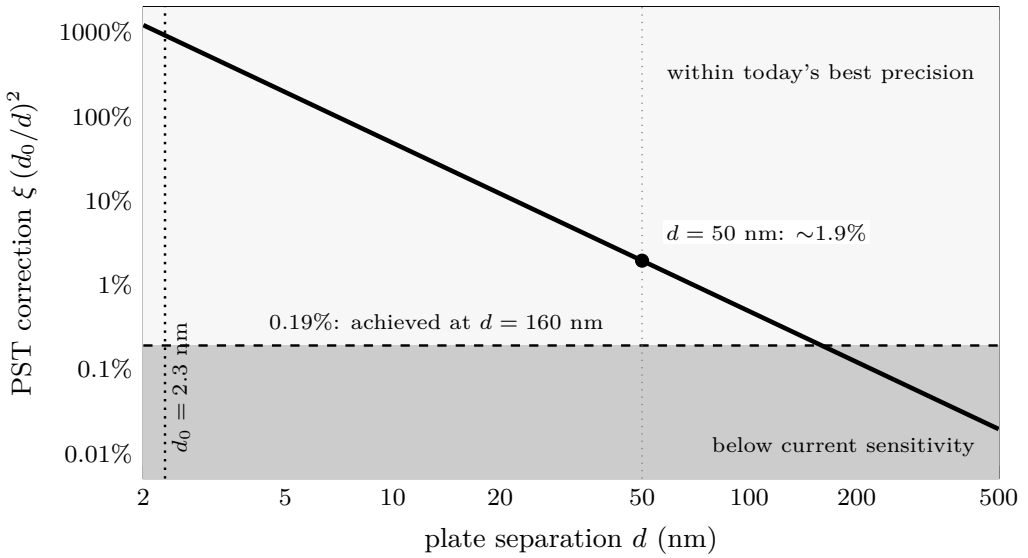


Figure 4: The PST Casimir correction $\xi\,(d_0/d)^2$ against plate separation, on logarithmic axes. The theory fixes the d^{-6} force law and the coefficient $\xi = 90/\pi^2 \approx 9.12$ exactly, so the *shape* of this curve carries no free parameter; only its amplitude scales with d_0, and $d_0 \lesssim 2.3$ nm is itself bounded by existing nanometre-Casimir data ((217), [212]). The shading marks experimental reach. Above the 0.19% line the correction is at or within the precision already achieved at 160 nm; below it, the correction is out of reach until that precision is carried to shorter separations. At $d = 160$ nm the predicted correction is 0.19%, sitting exactly at that floor by construction, since saturating it is what fixes the bound (eq. (209)); at $d = 50$ nm, reachable by the torsion-pendulum and atomic force microscope (AFM) experiments now under development, it rises to about 1.9% (eq. (210)). That is the sensitivity such a measurement would have, not the signal the theory expects. Note what a null result would mean: d_0 enters only as the amplitude, so measuring no deviation does not exclude PST, it lowers the bound on d_0 and pushes the curve down. What PST cannot survive is a correction with the wrong exponent. All of this is at the bound; under the canonical normalisation the curve sits eighteen orders lower (Computation 201).

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

$$\left(\frac{d_0}{d}\right)^2\Bigg|_{d=50\,\mathrm{nm}} = \left(\frac{2.3}{50}\right)^2 \approx 2.13 \times 10^{-3} \tag{210}$$

The correction with $\xi = 90/\pi^2$ is $9.12 \times 0.213\% \approx 1.9\%$, an order above the 0.19% floor already achieved at 160 nm; it would be detectable if that precision were carried down to 50 nm and d_0 sat at its bound, which under the canonical normalisation it does not.

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

$$\left(\frac{d_0}{d}\right)^2\Bigg|_{d=20\,\mathrm{nm}} = \left(\frac{2.3}{20}\right)^2 \approx 0.0133 \tag{211}$$

A correction of $9.12 \times 1.33\% \approx 12\%$ over standard quantum field theory (QFT) would be unambiguous at the bound: it would exceed 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, while none contribute at order d^{-6}, so the d^{-6} PST term is separable by fitting measurements at multiple separations and its ratio to the d^{-4} term grows as separation decreases (the categorical-exponent and two-layer-falsifiability structure is established in §15.6). Substituting the diagnostic upper bound $d_0 \approx 2.3$ nm explicitly into equation (208):

$$\delta P_{\mathrm{PST}}(d) = -\frac{\xi\,\pi^2\hbar c \times (2.3\ \mathrm{nm})^2}{240\,d^6} \tag{212}$$

The coefficient $\xi = 90/\pi^2 \approx 9.12$ is derived exactly from the Gaussian moments of the substrate weight function; see equation (213). Measuring the

d^{-6} component at two separations via the ratio test (218) would determine the amplitude scale d_0 from the data, with $\xi = 90/\pi^2$ fixed by the theory, and compare it with the value 2.2 am the normalisation premise fixes; measurements at additional separations overconstrain that one scale, providing an internal consistency check on the d^{-6} form.

15.6. The parameter-free exponent and falsifiability

A careful reader will observe that equation (208) 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:

$$\xi \;=\; \frac{15}{\pi^2}\,\frac{\displaystyle\int_0^\infty e^{-x^2/(4d_0^2)}\,x^4\,dx}{\displaystyle d_0^2\int_0^\infty e^{-x^2/(4d_0^2)}\,x^2\,dx} \;=\; \frac{15}{\pi^2}\times 6 \;=\; \frac{90}{\pi^2} \;\approx\; 9.12, \tag{213}$$

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

$$\int_0^\infty e^{-x^2/(4d_0^2)}\,x^4\,dx \;=\; 12\sqrt{\pi}\,d_0^5, \qquad \int_0^\infty e^{-x^2/(4d_0^2)}\,x^2\,dx \;=\; 2\sqrt{\pi}\,d_0^3. \tag{214}$$

Hence the numerator carries dimension d_0^5 (length^5) and the denominator carries $d_0^2 \times d_0^3 = d_0^5$ (length^5), 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 (213) 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 68, 73).

The coefficient $\xi = 90/\pi^2 \approx 9.12$ is therefore fixed exactly *given the Gaussian*

kernel choice; the amplitude $\xi\, d_0^2$ is parameter-free in its coefficient, with the scale d_0 the Landau–Ginzburg coherence length.

15.6.1. Kernel-shape sensitivity and the d_0 bound

The numerical value $\xi = 90/\pi^2$ is conditional on the Gaussian projection-kernel form. Computation 68 (Appendix B) tests the sensitivity of ξ across six plausible alternative kernel shapes, Gaussian, squared Lorentzian, single-sided exponential, hyperbolic-secant-squared, compact parabolic, and compact quartic-bump, by evaluating the moment ratio $M_4/(d_0^2 M_2)$ for each. Findings:

- *The d^{-6} power-law exponent is preserved across every kernel tested.* This is expected from the structural derivation above (parity plus the existence of a single length scale d_0), but Computation 68 confirms it numerically and isolates the kernel-dependence to the coefficient ξ alone.
- *The coefficient ξ varies by approximately a decade across the kernel family.* Concretely: $\xi_{\text{Gaussian}} = 9.12$; $\xi_{\text{exp}} = 18.2$; $\xi_{\text{sech}^2} = 21.0$; $\xi_{\text{compact quartic}} = 2.0$; $\xi_{\text{compact parabolic}} = 2.6$; $\xi_{\text{Lorentzian}} = 378.7$ (the Lorentzian outlier is from the $1/x^4$ heavy tail, which is physically implausible for a substrate discreteness scale). Restricting to the four kernels with exponential or compact-support tails, the range is $\xi \in [2, 21]$, a factor of ~ 10.
- *The diagnostic d_0 bound from current nanometre-Casimir data scales as $1/\sqrt{\xi}$.* Under the same data, the kernel-conditional bound ranges from $d_0 \lesssim 1.5$ nm (sech^2) to $d_0 \lesssim 4.9$ nm (compact quartic), with the Gaussian value $d_0 \lesssim 2.3$ nm between them. The order of magnitude (a few nm) is robust; the precise numerical bound is kernel-conditional with a factor-of-~ 3 spread.

The honest reading is therefore: *the d^{-6} signature is parameter-free and falsifiable; the amplitude ξd_0^2 that converts a measured deviation into an inferred d_0 carries a kernel-shape systematic of order a factor of two to three on d_0.* Reports of the bound should read "$d_0 \lesssim 1.5$–5 nm depending on the projection-kernel choice; the Gaussian (canonical) choice gives 2.3 nm." See

Computation 68 for the full kernel table.

15.6.2. The Gaussian as the substrate's kernel

The kernel-shape sensitivity above frames ξ as conditional on a *choice* of kernel. Computation 73 sharpens this by deriving the substrate's actual projection kernel directly from its Boolean structure and showing it is asymptotically Gaussian, with a calculable non-Gaussian correction that vanishes at the matched scaling. The substrate Hilbert space $\mathbb{C}^{2^D}$ carries the Boolean kernel $K_n(j) = \binom{n}{j}(1/2)^n$, the binomial distribution at $p = 1/2$ (the per-site Bernoulli probability of $\mu = \bigotimes_{a\in D} \mathrm{Bern}(1/2)$). By the central limit theorem (CLT), the binomial converges to a Gaussian as $|D| \to \infty$ at the matched scaling. The leading-order non-Gaussian correction is in the kurtosis:

$$\mathrm{kurt}(\mathrm{binomial}|_{p=1/2}) \;=\; 3 - \frac{2}{|D|}, \qquad \mathrm{kurt}(\mathrm{Gaussian}) \;=\; 3. \tag{215}$$

Substituting into the moment ratio M_4/M_2^2 that fixes ξ gives

$$\xi_{\mathrm{substrate}}(|D|) \;=\; \xi_{\mathrm{Gaussian}} \cdot \left(1 - \frac{2}{3\,|D|}\right) \;=\; \frac{90}{\pi^2}\left(1 - \frac{2}{3\,|D|}\right). \tag{216}$$

The kernel-shape conditional of Computation 68 is therefore not arbitrary: the substrate *is* approximately Gaussian by CLT, with the leading correction $-2/(3|D|)$ a small downward shift that vanishes at the matched scaling. At $|D| = 6$ the correction is 11%; at $|D| = 20$, 3.3%; at $|D| = 100$, less than 1%. The diagnostic $d_0 \lesssim 2.3$ nm bound derived from the Gaussian value is therefore correct to leading order; the substrate's actual bound is shifted by at most a few percent at any physically relevant $|D|$. This converts the kernel-shape conditional from "ξ depends on a choice of kernel" to "ξ is structurally Gaussian with a calculable convergence rate".

What does *not* depend on ξ or d_0 is the power-law *exponent*. As derived above for equation (208), the exponent -6 follows from the parity constraint and the existence of a single length scale d_0 alone, and uses neither the value of ξ nor of d_0. Odd powers of d_0/d would require a preferred substrate orientation,

inconsistent with the isotropy of Chapter 4. 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 one further quantity, d_0, the Landau–Ginzburg coherence length, which P1–P3 alone do not fix: it follows from the canonical normalisation of §12.2, the premise on which M_* rests, as $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am (Computation 201). Existing 160–750 nm precision Casimir data constrain d_0 only by an upper bound, $d_0 \lesssim 2.3$ nm (equation (217) at the derived $\xi = 90/\pi^2$), nine orders above that value. A measured residual at d^{-6} with amplitude A would therefore determine $d_0 = \sqrt{A/\xi}$ from the data, and any resolvable amplitude would put it far above 2.2 am and refute the normalisation rather than confirm the book. 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, and since the normalisation premise fixed d_0 (Computation 201) it is amplitude-rigid as well, on that premise: the exponent layer is parameter-free, and the amplitude layer is a prediction conditional on the canonical normalisation rather than a fit. The asymmetry that remains is between a structural consequence and a carried premise, and §12.2 names the premise.

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 §15.9. In brief: finite-conductivity corrections enter via the Lifshitz formula [17] 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.

15.7. Falsifiability and the experimental programme

The PST Casimir prediction satisfies strict experimental falsifiability, with the two-layer structure of §15.6: the d^{-6} exponent is the parameter-free content, while the amplitude carries the onset scale d_0, fixed by the normalisation premise at 2.2 am and bounded by the data at 2.3 nm. The prediction is refuted if:

1. A non-d^{-6} short-distance 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).

Null results, by contrast, tighten the bound on d_0 without refuting the structure, and they are what the theory expects at every reachable precision. Agreement with standard QFT to better than 1% at $d = 50$ nm would constrain $d_0 < 1.7$ nm, still nine orders above the derived value. Multi-separation ratio tests (equation (218)) in the 20–50 nm range, finding no d^{-6} component at the 0.5% level, would push the bound lower still. A d^{-6} component resolved at any reachable amplitude would confirm the form and refute the normalisation premise; the following would establish the form:

1. A d^{-6} component is detected at separations below 50 nm, with amplitude consistent with $d_0 \lesssim 2.3$ 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 that same d_0.
3. The fitted d_0 is stable across separations, apparatus, and materials: a single onset scale accounts for every ratio test, as a single Landau–Ginzburg coherence length requires.

Current experiments [212, 15] at 160–750 nm agree with standard QFT to 0.19% at 160 nm, consistent with $d_0 < 7.0\,\text{nm}/\sqrt{|\xi|}$ from equation (217), i.e. $d_0 \lesssim 2.3$ nm at the derived $\xi = 90/\pi^2$. The onset scale $d_0 = 2.2$ am the normalisation premise fixes lies nine orders below this bound, so current Casimir data are consistent with it and no planned experiment reaches the predicted amplitude; what the next generation of sub-50 nm measurements can do is lower the data's bound.

As established in §15.6, the d^{-6} exponent and the coefficient $\xi = 90/\pi^2$ are structural, while the onset scale d_0 is the Landau–Ginzburg coherence length. A measured d^{-6} residual at sub-50 nm separations would support the substrate-discreteness prediction, and a short-distance power law measured inconsistent with d^{-6} would falsify it. A null result, by contrast, only tightens

the data's bound on d_0, whose value the normalisation premise fixes far below it. The modal scale $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV is a separate, collider-testable prediction: PST expects new physics associated with the modal condensate near this scale. The substrate Walsh zero mode maps under the projection chain $\psi \to \Pi(\psi) \to H_{\text{doublet}}$ to the constant function on M (vacuum direction), structurally distinct from the emergent SM Higgs doublet generated by inner fluctuations on $M \times F$ (Computations 93 and 97). With that distinction explicit, the scalar-sector cancellation closes and $M_* = 4\pi m_h\sqrt{2/3}$ is a parameter-free one-loop derivation of the scalar sector on the canonical normalisation (§12.2); its full-SM extension is settled in the negative (Computation 120).

15.8. Observational constraint on d_0 from Casimir data

Precision Casimir measurements at separations 160–750 nm [212] report 0.19% agreement with standard QFT at 160 nm, at 95% confidence [212, 15]. Requiring $|\xi|(d_0/d)^2 < 1.9 \times 10^{-3}$ at $d = 160$ nm gives:

$$d_0 \;<\; \frac{160\text{ nm}\,\sqrt{1.9\times10^{-3}}}{\sqrt{|\xi|}} \;\approx\; \frac{7.0\text{ nm}}{\sqrt{|\xi|}} \tag{217}$$

With the PST-derived $\xi = 90/\pi^2 \approx 9.12$ (conditional on the Gaussian convolution-kernel form; §15.6), this constrains $d_0 \lesssim 2.3$ nm. The PST correction sits at that 0.19% precision for separations $d \lesssim 69\,d_0$ (where $|\xi|(d_0/d)^2 \geq 1.9 \times 10^{-3}$ at $\xi = 90/\pi^2$), which at the bound is $d \lesssim 160$ nm, the precision of existing measurements; it reaches the 1% level for $d \lesssim 30\,d_0 \approx 70$ nm. A ratio test between measurements at two separations d_1 and d_2 provides a clean, apparatus-independent signature:

$$\frac{P(d_1)/P_{\text{standard}}(d_1)}{P(d_2)/P_{\text{standard}}(d_2)} - 1 \;\propto\; \left(\frac{d_0}{d_1}\right)^2 - \left(\frac{d_0}{d_2}\right)^2 \tag{218}$$

This represents a falsifiable prediction distinguishable from both standard QFT and 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 (213)). The overall amplitude scales as $(d_0/d)^2$, with d_0 the Landau–Ginzburg coherence length of the modal condensate (§15.3). Its magnitude is fixed by the canonical normalisation at 2.2 am (Computation 201), seventeen orders above ℓ_P and nine below the diagnostic bound. The budget below is evaluated at the bound, $d_0 = 2.3$ nm, where the correction is at the edge of current precision near $d = 160$ nm, reaches several percent near $d = 50$ nm and 12 percent at $d = 20$ nm; that is the sensitivity a measurement would have, not the signal the theory expects, which at the derived d_0 is 10^{-19} of the force at 20 nm.

15.9. 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 $d_0 = 2.3$ nm (upper bound)	$\sim$ **1.9**%	$\sim$ **0.48**%
PST d^{-6} at $d_0 = 1.1$ nm (half the bound, illustrative)	$\sim 0.44\%$	$\sim 0.11\%$

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 the diagnostic upper bound $d_0 \approx 2.3$ nm, the correction is $\sim$1.9% at $d = 50$ nm, an order above the 0.19% precision already achieved at 160 nm and robustly detectable if that precision is carried down to 50 nm.

15.9.1. Predicted deviation across separations

The fractional deviation from the standard Casimir force is $\delta F/F = \xi\,(d_0/d)^2$ with the parameter-free coefficient $\xi = 90/\pi^2 \approx 9.12$; the amplitude-setting d_0 is 2.2 am under the normalisation premise and is bounded by the data at $d_0 \lesssim 2.3$ nm, the value the table is evaluated at. Table 1 gives the predicted deviation across the experimentally relevant separations and three representative coherence lengths. Boldface entries exceed the 0.19% floor achieved at 160 nm by the micromechanical torsional oscillator of [212] and are detectable if that precision is carried to the separation in question; the remainder are next-generation targets. Because the deviation scales as d^{-2} relative to the standard force while the systematics tabulated above carry distinct functional forms, the two-separation ratio test (equation (218)) isolates the d^{-6} component apparatus-independently.

15.10. 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. The full operator-by-operator Standard Model effective field theory (SMEFT) survey of empirical M_* bounds against LHC and electroweak precision observables (EWPO) is carried out in Computation 10, with the tree-level custodial SU(2) protection of ΔT confirmed at $\Delta T = 0$ (loop and new-states corrections to S and

Table 1: PST predicted fractional Casimir deviation $\delta F/F = \xi\,(d_0/d)^2$ (coefficient $\xi = 90/\pi^2 \approx 9.12$ parameter-free; amplitude set by d_0) versus plate separation d and three illustrative coherence lengths at the scale of the diagnostic bound; the theory's own $d_0 = 2.2$ am puts every entry at 10^{-19} or below. Boldface entries exceed the 0.19% precision achieved at 160 nm [212]. The dominant systematic (finite-conductivity Lifshitz) is d^{-1}-like and subtractable to $< 1\%$ from optical data; roughness ($\sim \sigma^2/d^2$) and thermal ($\sim k_B T d/\hbar c$) corrections sit at 0.01–1% with distinct d-dependences.

d (nm)	$d_0 = 0.5$ nm	$d_0 = 1.1$ nm	$d_0 = 2.3$ nm
20	**0.57**%	**2.8**%	**12**%
30	**0.25**%	**1.2**%	**5.4**%
50	0.09%	**0.44**%	**1.9**%
100	0.02%	0.11%	**0.48**%
160	0.01%	0.04%	0.19%

T are not computed there). 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 $\ell_* = \hbar c/M_* \approx 1.535 \times 10^{-19}$ 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 ($M_* \gtrsim 5$–8 TeV), SN1987A and neutron-star Kaluza–Klein (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 remained until Computation 205 was a genuine empirical-consistency question, not a derived constraint: whether $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV (a parameter-free one-loop derivation; Computations 93, 97) is consistent with the existing data once the PST-specific coupling structure (set by the projection operator Π) is computed. Computation 205 performs it under the

three readings this book's sentences admit of what M_* does to the fields below it. Read literally, as a momentum cutoff above which no field mode exists, it predicts no event with a parton momentum above 1.285 TeV; the ATLAS experiment measures dijet masses to 9 TeV and jet transverse momenta to 3.5 TeV with next-to-next-to-leading-order quantum chromodynamics (QCD) describing them [217], and holds 25 dilepton events above 2 TeV where the Standard Model expects 22 [218], so that reading is excluded outright. Read as the projection's Gaussian kernel on every vertex, it multiplies the Drell–Yan yield above 2.2 TeV by 0.05 against 19 observed dielectron events, and its linearised form is a contact interaction with an equivalent scale of 13 to 16 TeV against the published limits of 32 to 36 TeV [218], short by a factor five in $1/\Lambda^2$; a universal kernel would have to sit at $2.7\,M_*$ or above. Read as the kernel on the Higgs doublet alone, the field that *is* the projected condensate (§15.2), everything measured is consistent: the coupling modifiers shift by half a percent against $\kappa_V = 1.035 \pm 0.031$ and $\kappa_F = 0.95 \pm 0.05$ [219]; Higgs production above 450, 650 and 1000 GeV of transverse momentum is suppressed by 12, 23 and 45 percent against measurements whose uncertainties are of order the value [221]; Higgs pairs are suppressed by 27 percent against a limit of 2.4 times the Standard Model [220]; and because the whole doublet carries the kernel, Goldstones included, the cancellation that keeps longitudinal vector-boson scattering bounded is preserved, consistent with the measured same-sign WW rate [222], where a kernel on the physical h alone would break the partial-wave bound at 1.5 TeV. The surviving reading is the one §12.2 now states, and it predicts: a deficit in the Higgs transverse-momentum spectrum growing as $1 - \mathrm{e}^{-p_T^2/M_*^2}$, 23 percent above 650 GeV, within the reach of the high-luminosity run. The first measurement in that bin, ATLAS's boosted $H \to \tau\tau$ on Run 2 plus partial Run 3 [234], gives a Run 2 deficit above 650 GeV and a combined upper limit of 4.5 times the Standard Model at 95% confidence: a limit, not yet the measurement at 15 percent precision that would decide it, so the exposure stands untested. The coherence length d_0 that sets the absolute size of the Casimir signal is fixed by the canonical normalisation at 2.2 am (Computation 201), so the signal is predicted in form

(the d^{-6} scaling and the coefficient $\xi = 90/\pi^2$) and in size, far below present reach (§15.5).

15.11. Wilson coefficients: tree-level custodial ΔT protection

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, section 10 of the script).

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 $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$ 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 $\mathrm{SU}(2)_L$ 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 $c_{\mathrm{HWB}}(M_*) = c_{\mathrm{HD}}(M_*) = 0$, exactly.

Consequently the PST contribution to the Peskin–Takeuchi parameters is

$$\begin{aligned} \Delta S_{\mathrm{PST}} &= (16 \sin\theta_W \cos\theta_W / \alpha_{\mathrm{em}}) \cdot v^2/M_*^2 \cdot c_{\mathrm{HWB}}(M_*) = 0, \\ \Delta T_{\mathrm{PST}} &= 0. \end{aligned} \tag{219}$$

The SM-internal contributions to S and T (top quark, gauge-Higgs loops) are part of the SM reference prediction that the Particle Data Group (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 at the level of the established content (no new gauge bosons, no new light fermions, the substrate modes being higher harmonics of the Higgs). The hypothetical Boolean partner states that a full-SM boson–fermion cancellation at M_* would require (§12.2) are *not* part of this demonstrated content. That cancellation is settled in the negative (Computation 120), with §19.9.6 recording the closure, and the light-EFT conclusions here are unaffected. Four-fermion contact operators $(\bar{\psi}\gamma_\mu\psi)(\bar{\psi}\gamma^\mu\psi)$ require new vector or scalar mediators, which PST lacks, so their tree-level Wilson coefficients vanish. Read as a universal kernel on every vertex, the loop-suppressed values give effective scales $\Lambda_{\text{eff}} = M_*/\sqrt{c}$ of 13 to 16 TeV by chirality (Computation 205), below the current LHC Drell–Yan high-mass bound of 32 to 36 TeV [218], which is why that reading is excluded; in the reading of the cutoff that survives Computation 205 the gauge vertices carry no kernel at all, so four-fermion operators arise only through doublet exchange, Yukawa-suppressed, and sit far above the bound, while a kernel on the gauge vertices themselves is what the same data exclude. Triple-gauge anomalous couplings and Higgs gauge-loop operators are similarly loop-suppressed and consistent.

The astrophysical bounds (SN1987A neutron-star cooling, and large-extra-dimension graviton emission) do not transfer to PST, for the reason given in §15.10. PST has no Kaluza–Klein tower: the internal factor F is a finite noncommutative geometry with sub-fermi associated length, not a set of macroscopic extra dimensions.

The empirical consistency question raised earlier in this section is therefore settled by Computation 205 rather than by construction: $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV is not in tension with any existing collider or astrophysical bound in the Higgs-doublet reading of the cutoff, which is the one §15.2 states, and it is excluded by existing data in the two readings that would put the kernel on the emergent quanta.

15.12. Summary

This chapter 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 (and kernel-shape sensitive: a factor-of-$\sim$ 10 spread in ξ across plausible kernel families, Computation 68, becomes a factor-of-$\sim$ 3 spread in the bound inferred from it), 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, which the canonical normalisation fixes at 2.2 am (Computation 201); the electroweak modal scale is $M_* = 4\pi m_h\sqrt{2/3} \approx 1285$ GeV from m_h alone. At the diagnostic upper bound $d_0 \approx 2.3$ nm the correction would be $\sim$1.9% at 50 nm and $\sim$12% at 20 nm, within reach of next-generation precision Casimir experiments; at the derived d_0 it is 10^{-19} of the force at 20 nm, and a null result at every reachable precision is what the theory expects.

On the status of d_0: as grounded in §15.3, it is the Landau–Ginzburg coherence length of the modal condensate, $d_0 = \xi_\psi = \sqrt{c_{LG}/(2|a_{LG}|)}$, not the substrate's microscopic grain. Its magnitude is not an output of P1–P3 alone: with $|a_{LG}| = m_h^2/4$ it is the gradient normalisation in another variable, and the canonical normalisation on which M_* rests puts it at $\sqrt{2}\,\hbar c/m_h = 2.2$ am (Computation 201).

16. Cosmology: Acceleration from Ongoing Sublimation

The universe is not only getting bigger, it is getting bigger faster, and nobody is sure why. This chapter argues that if instantiation is still happening now, speeding up is exactly what you should expect, and then works out what believing that commits the theory to.

16.1. 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. Cosmic acceleration arises naturally from the thermodynamics of ongoing instantiation: conditional on the substrate's pre-spatial framing, Mechanism B of §9.1 delivers the equation of state $w = -1$ (the cosmological-constant form) via a steady-state-style non-dilution argument; Mechanism A separately gives the standard quantum-vacuum contribution at the Standard Model (SM) hierarchy scale.

16.2. The Cosmological Principle from the Bernoulli measure

The precausal configuration space is equipped with the Bernoulli product measure $\mu = \bigotimes_{a \in D} \mathrm{Bern}(1/2)$, equation (9), 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 substrate-internal saturation density of threshold-

crossed configurations, measured against the configuration-space measure $\mu = \bigotimes_{a \in D} \mathrm{Bern}(\frac{1}{2})$ of P3. Its dimension is events per substrate event-volume, consistent with the dimensional pinning in §16.5; it is *not* a per-proper-volume rate in emergent spacetime, which would yield Hoyle–Bondi–Gold steady-state $w = 0$. Its projection into the instantiated domain $\mathcal{G}$ is $\Pi(j_\tau)$. A spatial variation of the projected density $\Pi(j_\tau)(\mathbf{x}) \neq \Pi(j_\tau)(\mathbf{y})$ 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 $\Pi(j_\tau)$ would require a distinguished locus in the substrate, which μ forbids. Isotropy of $\Pi(j_\tau)$ 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 cosmic microwave background (CMB) uniformity [71], 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 [68, 179, 180]:

$$ds^2 = c^2\,dt^2 - a(t)^2 \left[\frac{dr^2}{1 - kr^2} + r^2\,d\Omega^2 \right], \tag{220}$$

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

16.3. 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 [52, 53] 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 [74, 71] is not evidence for inflation; it is direct observational evidence for the permutation invariance of the Bernoulli measure.

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

16.4. Ongoing instantiation: why it 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 spatially uniform by the Cosmological Principle theorem above and constant in time by non-dilution (below); the energy released per crossing event is $\Delta\mathcal{F} = \mathcal{F}_{\text{pre}} - \mathcal{F}\big|_{\mathcal{V}} > 0$, the descent from the high-tension pre-threshold configuration to the vacuum

manifold $\mathcal{V}$ (the sombrero valley of Chapter 4). The effective energy density of ongoing instantiation is therefore

$$\rho_{\text{inst}} = \frac{j_\tau \, \Delta\mathcal{F}}{c^2}. \tag{221}$$

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

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

Therefore:

$$\dot{\rho}_{\text{inst}} = 0. \tag{222}$$

16.5. 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,

$$\dot{\rho} + 3H(\rho + P) = 0, \tag{223}$$

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

$$P_{\text{inst}} = -\rho_{\text{inst}}, \qquad w \equiv \frac{P_{\text{inst}}}{\rho_{\text{inst}}} = -1. \tag{224}$$

This is the vacuum sector of (48), and it is the only one that can carry $w = -1$: Computation 155 shows that a stress-energy proportional to $g_{\mu\nu}$ requires exactly $\rho + p = 0$, and Computation 158 shows the matter sector

$\Pi_{\mu\nu}$ is traceless and so carries $w = 1/3$. The two add and do not compete (§7.4.4). 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

$$\Lambda_{\text{PST}} = \frac{8\pi G}{c^2}\,\rho_{\text{inst}} > 0. \tag{225}$$

This is positive, de Sitter, not anti-de Sitter. §9.1 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 contribution, positive at the naive cutoff but regulator-dependent in sign (*Scope of Mechanism A*, §9.1). The sombrero-valley energy $\mathcal{F}\big|_{\mathcal{V}} < 0$ is internal to the precausal substrate and does not couple directly to emergent gravity as Λ; the apparent AdS conclusion arises only from conflating $\mathcal{F}\big|_{\mathcal{V}}$ 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. 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 $j_\tau\,\Delta\mathcal{F}$. The Λ_{obs} *magnitude* (and its $\sim 10^{-120}$ hierarchy relative to M_{P}^2) is $T(C)$-contingent, via an order-of-magnitude analysis using the Cramér framework of Computation 75 (Computation 82): the magnitude depends on the explicit form of $T(C)$ as a functional of substrate configurations plus the threshold scaling $\tau(D)$, and both are configuration-content quantities by the structural-scope theorem of §14.3. The sign and equation of state $w = -1$ established above remain structural theorems (Computation 137).

16.6. Friedmann equations and cosmic acceleration

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

metric (220) yields the standard Friedmann equations:

$$H^2 = \frac{8\pi G}{3}\Big(\rho_m + \rho_r + \rho_{\text{inst}}\Big) - \frac{kc^2}{a^2}, \tag{226}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\Big(\rho_m + 2\rho_r - 2\rho_{\text{inst}}\Big), \tag{227}$$

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 (227) reduces to

$$\frac{\ddot{a}}{a} \;\approx\; \frac{8\pi G}{3}\,\rho_{\text{inst}} > 0, \tag{228}$$

giving late-time acceleration with no additional assumption. The role of ρ_{inst} in PST is structurally identical to Λ in ΛCDM, but the interpretation differs. In ΛCDM, Λ is a free parameter inserted by hand. Here, $\Lambda_{\text{PST}} = 8\pi G\,\rho_{\text{inst}}/c^2$ 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.

16.7. The rate problem

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

$$j_\tau = \int_{\mathcal{S}} \mathbf{1}\big[|T(C)| \geq \tau\big]\, d\mu(C), \tag{229}$$

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 Chapter 15. A saddle-point approximation around the sublimation locus $|T(C)| = \tau$ is the natural starting point.

The value of j_τ is not computed here, its magnitude being $T(C)$-contingent

like the Λ magnitude it fixes (Computation 82), but its structural form reduces (Computation 75) to a standard Bernoulli-measure large-deviation probability:

$$j_\tau \;=\; \mu\big(\{C \in \mathcal{P}(D) : |T(C)| \geq \tau\}\big), \tag{230}$$

which is the tail probability of $|T(C)|$ under the i.i.d. Bernoulli measure $\mu = \bigotimes_a \mathrm{Bern}(1/2)$ delivered structurally by P1. By Cramér's theorem [150], j_τ admits the asymptotic

$$j_\tau \;\sim\; \exp\big(-|D|\, I(\tau/\sqrt{|D|})\big) \quad \text{as } |D| \to \infty, \tag{231}$$

where I is the Cramér rate function, the Legendre transform [150] of the cumulant generating function of $|T(C)|/\sqrt{|D|}$. Computation 75 verifies this exponential-in-$|D|$ decay empirically on a representative random-walk-like $T(C)$ model. The remaining structural input from PST is the explicit form of $T(C)$ at the postulate level (beyond the random-walk-like model used in Computation 75) and the precise τ-scaling with $|D|$; both are $T(C)$-contingent by the structural-scope theorem of §14.3, and once both are supplied, Cramér delivers j_τ analytically. Thus j_τ is a Cramér saddle-point computation conditional on the explicit $T(C)$ functional, a structural reduction analogous to the A7 $\to$ uniform-rate reduction. The observed value $\Lambda_{\mathrm{obs}} \approx 1.1 \times 10^{-52}\ \mathrm{m}^{-2}$ provides a target: $j_\tau \cdot \Delta\mathcal{F} = \Lambda_{\mathrm{obs}}\, c^4/8\pi G \approx 5.3 \times 10^{-10}\ \mathrm{J\,m}^{-3}$. 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 Chapter 15.

16.8. Observational commitments

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

1. **CMB uniformity.** Permutation invariance of μ guarantees homogeneity and isotropy prior to any causal structure. Primordial fluctuations are Gaussian at leading order because the Bernoulli measure generates inde-

pendent increments; deviations from Gaussianity fall with the number of grains in a window, the skewness as $(d_0/\lambda)^{3/2}$ and the excess kurtosis as $(d_0/\lambda)^3$ for wavelength $\lambda \gg d_0$ (Computation 217). The spectrum at the grain is white; the amplitude and the tilt of the primordial spectrum are not derived (§19.9.4).

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 quantum mechanics (QM) sector, which emerges as a projection of the precausal structure (Chapter 5) and is indistinguishable from standard QM at energies well below $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV.

3. **Equation of state.** The instantiation component has $w = -1$ exactly at every epoch, so the instantiation-dominated epoch is de Sitter. This is the theory's sharpest cosmological commitment and its most imminent falsifier: it is already under test, and the present Dark Energy Spectroscopic Instrument (DESI) results favour a dynamical equation of state over ΛCDM at 2.8σ–4.2σ depending on the dataset combination [198]. What that would and would not cost PST, and why the theory cannot absorb it, is set out in §16.9.

4. **Spatial flatness.** The driving of $\Omega_k \to 0$ by ρ_{inst} and the absence of initial fine-tuning predict $\Omega_k = 0$ to high precision, consistent with current CMB bounds [71].

5. **Cosmic acceleration.** Equation (228) gives late-time acceleration consistent with Type Ia supernova data [69, 70] without introducing new fields beyond those required by the precausal structure.

16.9. What would refute the constancy, and what would not

Chapter 15 separates the Casimir prediction into a parameter-free *exponent* and an *amplitude* that rests on the normalisation premise, and insists that the exponent carries the structural falsifiability. The cosmological result has

the same two-layer structure, and it is worth stating it as plainly, because the layer that is exposed is not the one a reader is likely to expect.

The magnitude is not PST's to predict, and no measurement of it can refute the theory. The observed value of Λ and its $\sim 10^{-120}$ hierarchy are $T(C)$-contingent: they are properties of how this world came out, not consequences of P1–P3, and §16.7 says so without hedging. A revised measurement of the dark-energy density is not evidence for or against PST, in exactly the way a null Casimir result tightens the bound on d_0 rather than excluding the theory.

The constancy is PST's, and it is what a measurement can refute. §16.5 does not merely accommodate dark energy; it derives non-dilution, and non-dilution forces $w = -1$ *exactly*. The derivation fixes no amplitude and admits no tuning: $\rho_{\text{inst}} = j_\tau \, \Delta\mathcal{F}/c^2$ neither dilutes nor redshifts, so $\dot{\rho} = 0$ and continuity gives $w = -1$ at every epoch. This is a statement about *scaling*, in the same structural register as the d^{-6} exponent, and there is no contingent quantity behind which it can shelter.

16.9.1. The falsifier, stated sharply

The claim of §16.5 is refuted if:

1. The dark-energy equation of state is established to differ from -1 at any epoch in the instantiation-dominated era. PST predicts a constant, not a value near -1 today, so a confirmed $w_0 \neq -1$ and a confirmed $w_a \neq 0$ are each sufficient.
2. *A fortiori*, a confirmed crossing of the phantom divide. A history in which $w < -1$ in the past and $w > -1$ now is not a small departure from the prediction but is incompatible with the mechanism at every redshift: a density that does not dilute cannot cross. This matters because it is the shape the current data prefer.

The claim is *not* refuted by:

1. A measurement of Λ's magnitude, at any precision. That quantity is

$T(C)$-contingent and the theory never claimed it.

2. A failure to derive the 10^{-120} hierarchy. Same layer, same answer.
3. A tighter bound that remains consistent with $w = -1$, which confirms rather than threatens.
4. A statistical *preference* for a model carrying extra parameters, absent a determination that w departs from -1. Clause 1 above requires the departure to be established; a model comparison that favours a larger parameter space is not by itself that determination, particularly where the datasets entering it are in tension with one another.

The first three touch the contingent layer only. The fourth is a statement about evidential thresholds rather than about layers, and the next subsection returns to it, because it is the clause the current measurements actually engage.

A rival that predicts the departure. The nearest programme to PST on the cosmological constant is causal set theory's everpresent Λ [227]. There the discreteness of the causal set makes the four-volume V a counted quantity, Λ is its conjugate, and the Poisson fluctuation of the count gives $\Lambda \sim \pm 1/\sqrt{V}$: a cosmological "constant" of order the critical density at every epoch, wandering in magnitude and in sign. That is an answer to the why-now question which PST does not offer, since §16.7 leaves the magnitude contingent, and it is a prediction PST cannot share: a wandering Λ is a w that departs from -1, which is the first refuting clause above. Fits of everpresent Λ models reach the current data as well as ΛCDM does and ease the low-redshift tension in H_0 [228]. So the two programmes, which agree that the cosmological constant is a consequence of discreteness beneath spacetime, are separated by exactly the measurement this section names: a $w(z)$ established to be constant refutes everpresent Λ, and one established to wander refutes PST.

16.9.2. Where the measurements currently stand

On this question PST's prediction coincides with ΛCDM's: a cosmological constant, $w = -1$ at every epoch. The concordance model remains the one a

rival has to beat, and the position defended below is that PST's cosmological claim is at present *vulnerable and not refuted*; it has acquired a live route to refutation that it did not have when the derivation was written, and it has not been refuted. What follows is the evidence for that assessment, taken in the order a sceptic would want it: the challenge first, then its condition.

The Dark Energy Spectroscopic Instrument's second data release reports a dynamical dark-energy model preferred over ΛCDM at 3.1σ for DESI baryon-acoustic-oscillation data combined with the cosmic microwave background, rising to between 2.8σ and 4.2σ when Type Ia supernova compilations are added, the range reflecting which supernova sample is used [198]. The favoured region is the quadrant $w_0 > -1$, $w_a < 0$: phantom in the past, quintessence-like today, crossing $w = -1$ in between. That is clause 2 of the preceding subsection, and it is the most serious observational challenge the claim currently faces.

Four things bear on how much of a challenge it is, and the first comes from the measurement itself.

The baryon acoustic oscillation (BAO) data are well described by a cosmological constant. DESI's own analysis states that its results *are* well described by a flat ΛCDM model, while recording that the parameters its BAO data prefer sit in 2.3σ tension with those from the CMB [198]. The preference for evolution is therefore not something the BAO measurements show on their own terms; it emerges from combining datasets that do not agree with each other, and a two-parameter extension will relieve a tension between datasets whether or not the extension is physical.

On the supernova side the preference tracks a known low-redshift systematic. Correcting for systematics in the low-z ($z < 0.1$) sample reduces the preference for dynamical dark energy to below 2σ [201], and a critical review of the DESI-era evidence finds it sits at $\lesssim 3\sigma$ for most dataset combinations, reaching 3.8σ and 4.2σ only when the Union3 and DES-SN5Y supernova samples are included, tracing that enhancement to the same low-redshift subsample [199]. The

internal pattern is itself suggestive: adding Pantheon$^+$ *lowers* the significance to 2.8σ, below the 3.1σ obtained without supernovae at all, while the 4.2σ figure comes from the compilation in which the low-redshift problem has been identified.

On the BAO side the deviation is carried by particular tracers. The apparent departure from ΛCDM in DR2 is driven mainly by the LRG1 and LRG2 samples, and removing either restores concordance [202]. Separately, testing DESI DR2 against Pantheon$^+$ through the cosmic distance-duality relation, which follows from photon-number conservation in any metric theory and holds whatever dark energy is doing, reveals a redshift-dependent violation; accounting for it suppresses the preference for dynamical dark energy [200]. A mismatch of that shape can masquerade as evolution.

And the extension does not buy what it is often invoked for. Neither ΛCDM nor the two-parameter alternative reduces the sound horizon by the $\sim 7\%$ required to relieve the Hubble tension [199], so the dynamical model is not recommended by independent explanatory work it performs elsewhere.

Against all of which stands the strongest point on the other side, which should be stated rather than omitted: the trend is not an artefact of the two-parameter (w_0, w_a) form, since non-parametric reconstructions and alternative extensions recover similar dynamical behaviour. The argument therefore lives in the systematics of the datasets, not in the choice of parameterisation, and there it remains unsettled.

Two conclusions follow, and they must be held together. The preference strengthened rather than faded between the first and second data releases, and 3σ against a claim that admits no tuning is a warning, not noise. But the falsifier of the preceding subsection has *not* fired: clause 1 requires a departure from $w = -1$ to be *established*, and what exists is a preference, in some dataset combinations, whose significance falls below 2σ under one published systematic correction and dissolves under another. Preference is not establishment, and clause 4 of what would *not* refute the claim is precisely

this case.

The distinction matters here for the same reason it matters in Chapter 15: a theory that reads pressure as refutation abandons a result it is still entitled to, and one that reads it as noise learns nothing from it. The question will be settled within the decade by DESI's later data releases together with Euclid, the Rubin Observatory, the Roman Space Telescope and the Simons Observatory, and what is owed in the meantime is to have said in advance what each answer would mean.

16.9.3. Why PST cannot retreat, and what that costs

It could be supposed that $w = -1$ is stated conditionally on the substrate's pre-spatial framing, and that relaxing that framing would yield $w \neq -1$ while preserving the mechanism. It would not, and the reason is worth making explicit before the theory has any need of it.

The framing is load-bearing twice over. j_τ is a saturation density measured against the configuration-space measure of P3, and it is defined that way because the substrate has no spatial extent against which a proper-volume rate could be posed. Replacing it with a per-emergent-volume creation rate does not give a slightly different cosmological constant; it gives Hoyle–Bondi–Gold steady state, $w = 0$, which is excluded by an enormous margin. And the same choice is what forbids PST a spontaneous collapse rate (§10.3.4): a spontaneous localisation rate is a per-volume, per-time quantity of precisely the kind the framing excludes.

The equation of state and the collapse rule are therefore exposed together. Were an evolving dark-energy equation of state confirmed, PST would not lose the cosmological constant and keep the rest; it would lose the framing that also underwrites its account of measurement. Conversely, were spontaneous localisation observed in an isolated superposition, the damage would run the other way. Two results that look independent share one point of failure.

This is a worse position to be in than having two separate risks, and it is a better thing to have established now than to discover when the data demand

it. A derivation that finds its own flexibility only once it is needed is worth much less than one that mapped the flexibility in advance; the flexibility here is mapped, and there is none.

16.10. Comparison with steady-state cosmology

The continuous creation of energy density in Bondi–Gold–Hoyle steady-state cosmology [72, 73] is the closest historical precedent for PST's ongoing instantiation: both predict $\dot{\rho} = 0$ as a consequence of an ongoing creation process. The differences are decisive, and the equation of state is one of them: classic steady-state cosmology creates ordinary matter that dilutes, giving $w = 0$, whereas PST's non-diluting instantiation gives $w = -1$ (§16.5).

Mechanism. Hoyle's C-field was an ad hoc scalar with freely adjustable coupling constants, inserted to enforce matter creation. PST's ongoing instantiation needs no such insertion, and the reason is worth stating carefully, because the obvious version of it is weaker than the one available.

The obvious version appeals to size: the configuration space is infinite, so the supply cannot run out. That argument does not survive contact with μ, which is a *probability* measure and therefore carries total mass 1 however large D is. A process consuming measure at a fixed rate exhausts it in finite time, and an infinite index set does not change that.

The argument that does work is that there is nothing to consume. What §16.5 calls j_τ is a saturation *density* against μ, not a flux drawn from a stock, and the same section already excludes the alternative for an independent reason: a per-proper-volume creation rate yields the Hoyle–Bondi–Gold steady state $w = 0$ rather than $w = -1$. A draining reservoir is a per-volume flux of exactly that form, so the depleting reading is ruled out by the equation of state PST derives, before any question of supply arises (Computation 147). Crossing changes a configuration's modal status; it does not spend it. The reservoir has no bottom because it is not a reservoir.

What is created. Steady-state cosmology creates new matter inside the instan-

tiated 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 [74] was decisive. PST expects a relic background: the pre-threshold epoch is high-tension, and its relic is expressed as zero-point fluctuations of the instantiated field; the background's temperature is not derived here. The CMB is consistent with PST, not evidence against it.

Structure formation. Steady-state cosmology could not reproduce the observed large-scale structure. The Bernoulli measure gives Gaussian, scale-free fluctuations at the grain (Computation 217); the near-scale-invariant primordial spectrum Planck measures [71], $n_s = 0.965$, one power of k from white, needs the transfer from the grain to the horizon that this book does not derive (§19.9.4), with the transfer-function structure of ΛCDM below it.

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.

16.11. 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.

In terms of what is established: the Cosmological Principle, the homogeneity and isotropy of $\Pi(j_\tau)$, is a theorem (§16.2) following from the permutation invariance of the configuration-space measure μ, conditional on the substrate-internal j_τ definition. The equation-of-state value $w = -1$ is a structural theorem conditional only on the pre-spatial framing that P1–P3 carry (Mechanism B), following from $\dot{\rho}_{\text{inst}} = 0$ under that same j_τ definition, with the non-dilution conclusion a consequence of Π's frame-independence applied to the pre-spatial substrate rather than an added postulate. Two mechanisms are offered for the sign of the effect. Mechanism A is regulator-dependent: the bare Goldstone zero-point on $\mathcal{V} \cong S^1$, which inherits the cosmological-constant sign-fine-tuning problem. Mechanism B is structural: a positive ρ_{inst} from ongoing threshold-crossing, with the sign forced by $\Delta\mathcal{F} > 0$. The magnitude is not fixed here: rate-fixing to match the observed Λ is $T(C)$-contingent content (Computation 82), not an open derivation.

17. Standing Among Existing Frameworks

Plenty of people have argued that space, time or matter are not fundamental. This chapter sets PST beside them and says what it does that they do not.

17.1. 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 chapter relates PST to each major existing framework, identifying both what each correctly describes and precisely where PST deepens or extends it. The routine physics-framework comparisons are collected in the table below; the relationships that carry unique structural or genealogical content are then treated as prose.

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.

17.2. Comparison with established physics frameworks

Table 2 relates PST to the major established frameworks for gravitation, quantum theory, and quantum gravity. In each case the framework is correct within its domain; PST supplies the precausal foundation that the framework leaves unexplained.

Table 2: PST related to the major established physics frameworks.

Framework	**What it gets right**	**What it leaves open**	**How PST grounds it**
General Relativity [1, 10, 45]	Curvature by matter/energy; perihelion precession, light deflection, black holes, cosmic expansion, gravitational waves (LIGO), verified for over a century.	Why a smooth Lorentzian manifold exists and has that signature; stress-energy and G taken as inputs; geodesic choice an external initial condition; initial singularity [2, 3]; diffeomorphism invariance postulated.	General relativity (GR) emerges as the conservation law $G_{\mu\nu} = 8\pi G\, T_{\mu\nu}$ with $T_{\mu\nu} = \Lambda\, g_{\mu\nu} + T_{\mu\nu}^{(\text{fields})}$, the vacuum term plus the projected matter fields: total asymmetric tension preserved across the threshold as curvature and stress-energy. Diffeomorphism invariance is the non-uniqueness of ρ; G is a unit conversion factor of Π; the geodesic gap closes via the sombrero topology of the slice vacuum $\mathcal{V}_{\bar{\eta}} \cong S^1$; the initial singularity dissolves because loci are not primitive.
Quantum mechanics (QM) / quantum field theory (QFT) [23, 24, 25]	All non-gravitational physics to one part in 10^{12} (electron $g-2$, Lamb shift); Hilbert space, Born rule, superposition, uncertainty, operator algebra from a compact formalism.	Hilbert space, complex amplitudes, Born rule and the collapse postulate assumed; QFT needs a fixed background; the QFT vacuum energy exceeds the observed Λ by ~ 120 orders of magnitude.	QM is the near-threshold regime $\varepsilon(C) \approx 0$: superposition is genuine intermediacy between wells, collapse the settling of ψ into one minimum. QFT is the fluctuation spectrum about $\mathcal{V}$ promoted by Φ (radial $m_h^2 = 4\varepsilon$, phase massless). Born rule (§10.2), collapse (§10.3), and spin-statistics (§11.3) are derived; the interacting vacuum exists nonperturbatively in the regulated scalar sector (§12.2, Computation 200), the gauge and fermion sector being the cutoff Standard Model's own; the vacuum energy is *relocated* to $\mathcal{V}$, its finite value $T(C)$-contingent (Chapter 16).

continued on next page

Table 2 continued from previous page

Framework	What it gets right	What it leaves open	How PST grounds it
Thermodynamics & black-hole entropy [46, 47, 19, 18]	Horizon entropy $S = A/(4\ell_P^2)$ and temperature; Jacobson derives the field equations from local horizon thermodynamics; Verlinde's entropic-force extension.	Why geometry, thermodynamics, and information connect is not explained by GR or QM alone.	*The thermodynamic half only.* $S = A/4G$ follows from the projected Einstein–Hilbert Lagrangian by Wald's Noether-charge theorem, which contains no microstate count, once the field equations are shown to reach the horizon; they do, by some forty-two orders (§8.3.3, Computations 154 and 216). A *statistical* entropy would instead count the modal degeneracy (configurations C projecting via Φ to the same geometry, ρ being non-injective) and PST supplies no such count (§19.9.3). ρ's fibres *are* the spacetime points, so its non-injectivity is generic rather than horizon-specific, and Π's kernel $\delta(x-\rho(C'))$ is supported on the bulk, not on a boundary. What PST does offer here is a precausal reading of the Hawking temperature, and the remark that Jacobson's derivation is the entropic reformulation of (48).

continued on next page

Table 2 continued from previous page

Framework	What it gets right	What it leaves open	How PST grounds it
Causal Set Theory [4, 5, 232]	Background-independent discretisation whose Poisson sprinkling is provably Lorentz invariant [225]; order plus number recovers the continuum metric by Malament's theorem [223]; scalar fields carried on the order [226]; a fluctuating cosmological constant, everpresent Λ [227].	Takes the causal partial order $\leq$ as primitive, the very structure that needs grounding; almost every finite order is not manifold-like [224], so a dynamics must select the ones that are; gauge structure and fermions are not derived from the order.	Both metric and causal order are derived; the causal order has no pre-image in $\mathbf{S}$ and is generated by the Lorentzian signature from Φ (Chapter 5). Discreteness arises near threshold where ρ is sensitive to pairwise tensions; at large $\varepsilon(C)$ the continuum is recovered. The two part on three points that can be measured or proved: PST predicts $w = -1$ exactly where everpresent Λ predicts a w that wanders (§16.9); PST's dimension falls to four from above, the direction the full-causal-set diffusion also reports (§20.6, Computation 207); and the sprinkling theorem, that no finite-valency graph is Lorentz invariant, which PST's cloud passes exactly when its realised positions are independent, a condition A2 does not fix and which the maximum-entropy clause supplies, read as a law on the realised cloud: the conditioned cloud is Poisson at the grain in every frame and hyperuniform only at the Mosco scale (Computation 208, Computation 211; the explicit scaling placement cannot test it, Computation 209; Appendix D).
Loop Quantum Gravity [6, 7]	Background-free quantisation of geometry; discrete area/volume spectra; spin networks with SU(2) labels; a loop-cosmology bounce.	Barbero–Immirzi γ fixed only by matching BH entropy; semiclassical limit incomplete; the problem of time.	Spin networks describe near-threshold non-uniform tension; SU(2) labels encode $L \in \hbar\mathbb{Z}$ from the U(1) vacuum symmetry; the problem of time dissolves since time is coinstantiated through Φ, so the Hamiltonian constraint is a consistency condition on the projection.

continued on next page

Table 2 continued from previous page

Framework	What it gets right	What it leaves open	How PST grounds it
String / M-Theory [26, 27]	Graviton as a massless spin-2 mode; finite amplitudes at all orders; gauge bosons and fermions in the superstring spectrum; mirror symmetry, D-branes, AdS/CFT (anti-de Sitter/conformal field theory).	Ten (eleven) dimensions imposed for consistency; six compactified on a Calabi–Yau [166] not fixed by the theory; $\sim 10^{500}$-vacuum landscape with anthropic selection; still needs a background; why strings is unanswered.	PST's "ten" is the total real K-theory (KO) dimension of a real spectral triple (modulo eight by Bott periodicity), $4 + 6$, with a finite noncommutative internal factor $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ of no spatial extent: no Calabi–Yau, no worldsheet, no landscape. Ten is derived as the unique minimum total KO-dimension compatible with a Lorentzian fermion action; C is a specific point in the modal potential, not a choice from a large set. (Scope contrast as prose below.)
Kaluza–Klein [28, 29]	Four-dimensional gravity plus electromagnetism from five-dimensional geometry; the U(1) of electromagnetism as the isometry of a compactified circle.	The extra dimension and its Planck-scale compactification radius are imposed by hand.	The Kaluza–Klein circle is the vacuum manifold $\mathcal{V} \cong S^1$ of the sombrero potential, a modal (phase) degree of freedom of the complex order parameter, not an independent spatial direction; the conserved U(1) charge is that of the broken neutral direction the Z absorbs, not electromagnetism (§8.2.1). No compactification by hand is needed.
Large extra dimensions / braneworld [16, 30]	Extra-dimensional effects at short distances/high energies; Arkani-Hamed–Dimopoulos–Dvali (ADD) lowers the gravity scale to TeV; Randall–Sundrum warps a fifth dimension to address the hierarchy.	Number and size of extra dimensions and the bulk geometry are inputs; why extra dimensions exist is not answered.	PST has *no* extra spatial dimensions: the internal factor is a finite KO-dimension-six noncommutative algebra, and d_0 is the Landau–Ginzburg (Casimir) coherence scale, not a compactification radius, so collider/astrophysical bounds do not constrain it *a fortiori*. The substantive content, the d^{-6} Casimir power law, is structural (kernel-independent), with coefficient $\xi = 90/\pi^2$ under the Gaussian projection kernel and amplitude set by d_0. The early high-dimensional geometry is a near-threshold consequence, reducing to four as instantiation grows.

continued on next page

Table 2 continued from previous page

Framework	What it gets right	What it leaves open	How PST grounds it
Causal Dynamical Triangulations (CDT) [31]	Non-perturbative sum over causal geometries; spectral dimension flowing from two at the Planck scale to four at large scales, from first principles.	The dimensional flow is a quantum average over triangulations with causality imposed by hand at the level of histories.	PST predicts the same endpoint by an explicit mechanism and the *opposite* flow: $r_0 = \sqrt{\varepsilon/(2b)}$ grows with $\varepsilon(C)$, the geometry exploring fewer directions until $\mathcal{V}$ settles it at four, so the dimension is high near threshold and falls to four where CDT's rises from two (§20.6; the full-causal-set diffusion shares PST's direction [229], and Computation 207 computes the spectral dimension on the i.i.d. surrogate: above four between the grain and d_0, four above it, two nowhere, the rise below d_0 seen by a probe of the substrate's walk rather than of the emergent geometry, Computation 210). No sum over geometries and no imposed causal constraint, since causality itself is derived; PST supplies CDT's precausal origin of that constraint.
Holography / AdS-CFT [32, 33]	Information bounded by boundary area, $S = A/(4\ell_P^2)$; AdS/CFT as an exact bulk-boundary duality.	The holographic bound is taken as a principle rather than derived.	Holographic encoding is the precausal reduction of effective dimensionality as geometry grows: $\rho : D \to M$ projects fewer independent directions as the order parameter settles, compressing the bulk onto the boundary. The bound is *not* derived: Π's kernel is supported on the bulk, and no area law follows from it (§19.9.3). AdS/CFT is this mechanism in an anti-de Sitter bulk with a specific asymptotic structure.

17.3. String theory: the scope contrast

PST and string theory both arrive at the number ten but for unrelated reasons, and the coincidence should not be mistaken for kinship; the structural distinction (string theory's compactified target manifold versus PST's finite noncommutative KO-dimension) is summarised in Table 2. The deeper contrast

is at the level of derivation-claims, where the two take opposite postures towards the structure / configuration distinction. String theory aspires to derive "everything from nothing": a single fundamental string theory should fix all coupling constants, all masses, and the gauge structure with no further input. The string landscape is the price of the aspiration: the equations admit $\sim 10^{500}$ vacua, each delivering a different low-energy physics, and no principle inside the theory selects the observed one. The honest post-landscape position is that string theory does not, in its present state, predict the specific physics observed, it accommodates a large class of possibilities, one of which is the Standard Model. PST takes the opposite posture. It identifies, structurally and explicitly, which quantities live in the postulate-derived layer: the gauge group, the generation count, the Higgs mass relation (scalar-sector; §12.2, Computation 120), spacetime emergence, the d^{-6} Casimir power law, $w = -1$ for the cosmological term, and so on. And it identifies which live in the configuration content $T(C)$, as boundary data of *this* universe's substrate: the Yukawa hierarchy, Cabibbo–Kobayashi–Maskawa [CKM] mixing, the cosmological-constant magnitude, and the analogous neutrino-sector content; the coherence length d_0 is not on this list, being fixed by the normalisation premise instead (Computation 201). This honest scope identification, "these specific items live in $T(C)$", is more rigorous, and more predictive, than the alternative of claiming to derive everything while in practice delivering a landscape with no specific predictions. A foundational theory's job is to make the structure / configuration distinction precise; PST does so explicitly, and the resulting contingent set is four families of items rather than the Standard Model's 19+ free parameters.

17.4. 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.

PST also extends the Chamseddine–Connes–Marcolli construction. Because the standard heat-kernel spectral-action machinery requires a continuous spectrum and is structurally obstructed for the discrete substrate triple at matched scaling, PST replaces it with a partition-function-level analogue operating on substrate observables under the Bernoulli measure; this delivers the substrate-side factor $Z^2 = \mathrm{e}^{-1}$ of the Higgs quartic at the matched scale (one-line result; full account in §19.7).

17.5. Other approaches to the Higgs mechanism

The discovery of the Higgs boson at the Large Hadron Collider (LHC) [87, 88] confirmed a fundamental scalar field with non-zero vacuum expectation value, completing the verification of the Standard Model. The standard mechanism, electroweak spontaneous symmetry breaking (SSB) by a fundamental scalar

[75, 76, 77, 78, 79], nonetheless leaves several foundational questions open: why a fundamental scalar exists, why its potential is sombrero-shaped, why the vev is 246 GeV and the mass 125 GeV, and why the electroweak/cutoff hierarchy is what it is. PST's central claim (§15.2) is that the Standard Model (SM) Higgs is not an additional postulate but the four-dimensional projection of the substrate's modal order parameter, $\phi_{\text{Higgs}} = \Pi(\psi)$: the Higgs and the modal condensate are the same field at different scales. PST inherits the SM mechanism at the projected level and answers those "why" questions. The scalar is the projection of ψ; the sombrero shape is the Landau–Ginzburg (LG) structure of threshold P3; the vev is the radius r_0 of the vacuum manifold; and the mass is the radial LG mode $m_h^2 = 4\varepsilon$ (equation (122)). M_* itself is fixed by the scalar self-loop matched-scale relation $m_h^2 = (3/2)M_*^2/(16\pi^2)$ (§12.2), and the hierarchy problem is softened to a residual fine-tuning $\Delta \approx 2.3$ (§12.2) rather than the $\sim 10^{32}$-order SM tuning. PST relates to the principal alternatives as follows.

Connes–Chamseddine–Marcolli spectral action [48, 49]: the Higgs arises as an inner fluctuation of the Dirac operator on $M \times F$ with $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$. PST extends this in the foundational direction: the spectral triple itself is derived from P1–P3 under A1–A7 via Mosco convergence and Dixon's hyperspinor synthesis (§13.7, §18.3). At the matched scale M_* it also replaces the heat-kernel correspondence with a partition-function-level analogue yielding $\lambda_{\text{SM}}(M_*) = b\,\mathrm{e}^{-1} \approx 0.092$ (full account in §19.7).

Technicolor and composite Higgs [89, 90]: the scalar is a bound state of new strong fermions. PST disagrees: ψ is a scalar map on configuration space derived from the asymmetric tension of P2 (Chapter 3), not a composite; PST predicts no compositeness signatures or high-momentum Higgs form factor.

Coleman–Weinberg radiative SSB [91]: SSB driven by radiative corrections. In PST the threshold τ is structural, set by the substrate's directional bias when $|T(C)|$ exceeds τ, not by loops; Coleman–Weinberg corrections enter only as a projected-QFT refinement (§12.2).

Little Higgs / pseudo-Goldstone Higgs [92]: a light Higgs from collective symmetry breaking. PST partially aligns: the three angular dofs are Goldstones eaten by $W^{\pm}$, Z (§13.4), but the physical radial mode is the LG amplitude, its lightness set by the cutoff $M_* = 4\pi m_h\sqrt{2/3} \simeq 1.285$ TeV (§12.2), not collective symmetry.

Supersymmetric / extended sectors (the Minimal and Next-to-Minimal Supersymmetric Standard Models, MSSM and NMSSM): multiple doublets with superpartners stabilising the mass. PST predicts only the single Standard-Model doublet $H \in \mathbb{C}^2$, with M_* addressing naturalness without supersymmetric cancellations; no second doublet, charged Higgs, or supersymmetry (SUSY) partners, so LHC null results are consistent with PST and constrain those sectors against it.

17.6. Inflationary Cosmology and the Multiverse

Inflationary cosmology [52, 53] 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. A scale-free spectrum at the grain follows from the Bernoulli measure (Computation 217); whether the near-threshold fluctuation structure of $\delta g_{\mu\nu}$ in equation (109) carries it to the near-scale-invariant primordial spectrum, with its measured tilt, is not derived here (§19.9.4). 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 the observed one is selected anthropically; there is one precausal substrate whose single continuous instantiation includes everything that is.

That claim can be made sharper, because P3 appears at first to push the other way: crossing is necessary rather than optional, many configurations lie above the threshold, and nothing so far has said why they do not sublimate into geometries of their own. The answer is that every route to a second instantiation moves along a symmetry the theory already carries, so the multiplicity is gauge rather than physical (Computation 149). A different threshold direction $\hat{\tau}' \in S^6$ is a G_2 rotation, and G_2 is transitive on S^6, so the orbit space is a single point and $\hat{\tau}'$ yields a rotated copy of the same Φ rather than a second world. A different realisation map ρ is not an alternative universe either: that non-uniqueness *is* the diffeomorphism invariance of §4.12, and diffeomorphic spacetimes are one spacetime. And a disjoint region of the property domain does not detach, because the Bernoulli measure factors over any bipartition $D = D_A \sqcup D_B$ and the Hilbert space factors with it

(Computation 142), giving subsystems of one state space rather than two disconnected ones. Quotienting by these three leaves exactly one element. This is the argument §3.4 already uses to empty the residual Z_2 tie of content: alternatives related by a symmetry of everything are not alternatives.

The scope is worth stating precisely, since the multiverse question attracts overclaiming from every direction. What is settled is structural: nothing in P1–P3 generates a second, causally disconnected geometry. That is not a proof that no such thing could exist for reasons outside the postulates, and it is not a verdict on the eternal-inflation multiverse, which is a claim about a different mechanism in a different framework and is treated on its own terms above.

17.7. Spencer-Brown and Distinction-Based Ontologies

G. Spencer-Brown's *Laws of Form* [61] 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* [67] provides an earlier philosophical foreshadowing. The opening dialectic moves from pure being, which has no determinate content, to nothing, and then resolves both in becoming, the first concrete category. The precausal substrate has no geometric or causal content (the analogue of pure being collapsing towards nothing), while modal sublimation is the tenseless relation in which indeterminate tension stands to 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.

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

John Wheeler's "it from bit" thesis [62] is the most influential information-theoretic predecessor to PST. Wheeler's central claim is that every particle, field, and spacetime event derives from binary distinctions, from answers to yes-or-no questions registered by physical apparatus. The universe is participatory: it comes into being through the act of observation.

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

Tegmark's Mathematical Universe Hypothesis [63] 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.

17.9. Structural Realism and Bohm's Implicate Order

Ontic structural realism (OSR), as developed by Ladyman [64] and French [65], 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 [66] provides a complementary comparison. Bohm proposed that the apparent discreteness and separateness of physical objects constitute an explicate order projected from a deeper implicate order in which everything is enfolded into everything else. The holistic character of the implicate order and its projection into local explicate structure closely parallels the PST functor $\Phi : \mathbf{S} \to \mathbf{G}$, which projects the holistic precausal substrate into the locally structured instantiated geometry. The key difference is formalisation: Bohm's implicate order is developed at the level of physical intuition and interpretive framework without a mathematical axiomatisation, and no mechanism specifies when or how the projection occurs. PST provides the structure that Bohm's picture envisions: the precausal substrate is the implicate order, and Φ is the explicate projection, with the bifurcation condition $|T(C)| > \tau$ specifying precisely when projection occurs.

17.10. How the claims in this book are meant to be checked

The frameworks compared in this chapter disagree with one another about almost everything, and they have one property in common that matters more than any of the disagreements: each can be argued with. A framework that cannot be argued with is not a competitor to them whatever it resembles on the page, and the distinction is procedural rather than rhetorical. This section states what this work does so that its claims can be located, tested and, where wrong, shown to be wrong. It is a description of method, not a claim of merit.

1. *Every claim is graded, in one place.* Appendix C sorts the results into postulates, modelling choices, derived structural theorems, theorems that hold under a modelling choice, conditional theorems with their assumptions listed, identifications, and $T(C)$-contingent content that the postulates do not fix. A reader can determine which register a statement belongs to without reconstructing the argument that produced it, and the register is

not flattering: a result that is conditional is filed as conditional.

2. *Every computation is published and can be run.* The 217 scripts of Appendix B are available at the address given in the front matter, each tied to the section it supports, and every claim resting on a computation names that computation rather than asserting the outcome. A disputed numerical result can be re-executed rather than argued about.

3. *Premises are named where they are not proved.* Bridge premise (B) and framework premise F4 are carried as explicit unproved content, not folded into a derivation and left for a reader to discover, and the two are not independent: (B)'s remaining content *is* F4 plus an empirical renormalisation-group test. *Faithfulness* (6), P1(iii), is carried here: it is an ontological clause of P1 rather than a fourth postulate, the count of postulates still being three, but it is not derivable from (4) and (5) and Computation 184 exhibits the counterexample. *The non-balance condition* $\left|\sum_{a\in D} v(a)\right| > 0$ is carried here too. It was taken off this list at Computation 169 and is restored: the argument that removed it priced a distinguished direction where the condition needs a correlation priced, and Computation 152's own counterexample refutes it (§3.5). One item did come off and stays off, the extremal selection principle at Computation 168, which rests on reading P1 as supplying no selector, a reading this book uses throughout and states as one. Where a reduction closes part of a premise and leaves the rest, the text says which part.

4. *Predictions carry stated falsifiers and stated thresholds.* Chapter 15 gives the Casimir prediction a falsifier in its exponent and says explicitly that a null result is not one; §16.9 does the same for the equation of state, including what evidence would *not* refute it; §10.3.4 does it for the collapse rule. In each case the refuting observation is named in advance, and so is the standard it must meet.

5. *Open problems are listed rather than buried.* Chapter 19 separates what has been resolved from what remains open and states the second at the

same level of detail as the first, including items where the honest answer is that no derivation exists.

None of this makes the results correct. A theory can be scrupulously documented and wrong, and the tests named in item 4 are what decide the matter. What the five properties buy is something narrower and more useful: they make error *locatable*. A reader who believes a conclusion here is mistaken can name the computation that fails, the premise that is doing unadvertised work, or the clause whose threshold has been met, and say why. That is the only arrangement under which disagreement is productive rather than merely opposed, and it is the standard against which this work asks to be read.

17.11. Summary

This chapter places PST in the landscape of existing physical and philosophical theories. General Relativity is recovered on the semidefinite metric-traceless cone that the tension projection reaches, given the three postulates of §1.7 under A1–A7 (Chapter 18): 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. The chapter is comparative and contextual rather than a source of new claims: its purpose is to identify what each existing framework (general relativity [GR], quantum mechanics [QM], string theory, loop quantum gravity [LQG], causal sets, noncommutative geometry [NCG], technicolor, Coleman–Weinberg SSB, Little Higgs, SUSY) shares with PST and where PST departs. Substantive technical claims about PST itself are made in the preceding chapters; this chapter is principally comparative, the exception being the gauge-quotient argument of §17.6 (Computation 149) and the partition-function reading of the spectral action.

18. Convergence of the Projected Geometry

This is the hardest mathematics in the book, and it sits this far back on purpose. It proves that the jump from a grainy substrate to smooth space really works. The rest of the book needs that proof to exist; no reader needs to follow it to understand the physics.

18.1. Mosco convergence: theorem and assumptions

The Mosco-convergence statement used above (§7.8) is made precise here: the seven assumptions A1–A7 below are stated explicitly, the structural ones are verified, and the equidistribution condition A3 is supplied by P3's maximum-entropy sublimation. Computation 4 supplies finite-$|D|$ numerical verification of the spin-structure component. A fully rigorous proof of the full Mosco theorem with explicit functional-analytic estimates would expand this section to a companion paper of its own (cf. [98, 113] for the analogous arguments in Riemannian settings). §18.2 supplies the proof-detail core of that programme: it proves the load-bearing K-uniform Berry–Esseen rate, the single analytic estimate to which A7 reduces, for the i.i.d. surrogate, via the five-step Berry–Esseen argument and the three closures (a)–(c). Thus the present subsection states the theorem and the reduction of A7 to that rate, and §18.2 proves it for the surrogate at proof-detail level, corroborating the rate that P3's maximum-entropy sublimation supplies for the emergence-determined case. The present form makes precise what PST claims, what is structural, and what P3 supplies.

18.1.1. Setup: the substrate sequence and matched scaling

Let D_n be a sequence of substrate sites with $|D_n| \to \infty$ and matched scaling parameter a_n (the substrate site spacing, distinct from the fixed coherence length d_0 of Chapter 15 and of the A4 displacement covariance) such that $|D_n| \cdot a_n^4 \to V_M$ (the Riemannian volume of M, read locally as in A1 below) and $a_n \to 0$ (the standard quantum Gromov–Hausdorff (QGH) convergence scaling [98]; cf. Computation 62). Define the configuration-space probability

space $(\mathcal{P}(D_n), \mu_n)$ with $\mu_n = \bigotimes_{a \in D_n} \mathrm{Bern}(\frac{1}{2})$; the Boolean Dirichlet form $E_n^B : L^2(\mathcal{P}(D_n), \mu_n) \to [0, \infty]$ as in eq. (55); the realisation map $\rho_n : D_n \to M = \mathbb{R} \times S^3$ with image $\rho_n(D_n)$ asymptotically equidistributed with respect to the Riemannian volume of M; the pullback $\Pi_n^\dagger : C_c^\infty(M) \to L^2(\mathcal{P}(D_n), \mu_n)$ via the realisation map; and the target Dirichlet form $E^{\mathrm{Riem}} : H^1(M) \to [0, \infty]$ generating the Laplace–Beltrami operator Δ_g.

18.1.2. The seven assumptions A1–A7

The structural and analytic inputs needed for Mosco convergence collect into seven explicit assumptions.

A1 (Matched scaling). The sequence (D_n, a_n), with site spacing $a_n \to 0$, satisfies $|D_n| \cdot a_n^4 \to V_M$, the Riemannian volume of M. This is the QGH-matching condition (cf. Computation 62) and is an input from the PST architecture, not a substrate consequence. (Since $M = \mathbb{R} \times S^3$ is non-compact, V_M is infinite, so A1 and A2 are read locally. The local node density is $|D_n \cap K|/\mathrm{vol}(K) = \lambda_n = 1/a_n^4 \to \infty$ on every compact K, the sites densifying as $a_n \to 0$. That is all the rate estimates require, being carried out on compact K with $u \in C_c^\infty(K)$.)

A2 (Equidistribution of the realisation map). For every $f \in C_c^\infty(M)$,

$$\frac{1}{|D_n|} \sum_{a \in D_n} f(\rho_n(a)) \;\longrightarrow\; \frac{1}{V_M} \int_M f \sqrt{-g}\, d^4x. \tag{232}$$

Verified for the i.i.d. surrogate by the strong law of large numbers. For the emergence-determined ρ_n the substrate's permutation invariance of μ delivers equidistribution *in measure*; the quantitative weak (H^1-tested) discrepancy rate the convergence argument uses is then exactly the rate condition A3, supplied by P3's maximum-entropy sublimation (§18.1.6).

A3 (Equidistribution condition, supplied by P3). For the emergence-determined cloud ρ_n, the metric-covariant unevenness of the realised node measure against $\sqrt{-g}\,\mathrm{d}^4x$ stays within the closure margin, i.e. the rate hypothesis of §18.1.6 holds with some $\eta_{K,R} < d_0^2/4$. This is supplied by P3's

maximum-entropy sublimation, not a separately adopted premise. A2 establishes equidistribution *in measure* (forced by the permutation symmetry of μ); maximum-entropy sublimation strengthens this to the realised configuration itself, selecting the μ-typical instantiated configuration, the analogue of smoothness (or generic initial data) in general relativity and of typicality in statistical mechanics. It is a typicality principle internal to P3, not a fourth postulate and not an analytic derivation of the rate from more primitive structure. The i.i.d. surrogate computations (Computations 70, 83) corroborate the rate P3 supplies. The same clause supplies one thing more: that the realised positions be independent, Poisson-type, which equidistribution in measure does not fix (a lattice satisfies A2 and selects a frame) and which is the hypothesis under which the sprinkling theorem of causal set theory makes PST's discreteness Lorentz invariant (Computation 208). Read as a law on the realised cloud, maximum entropy given the density is independence, by the subadditivity of entropy, and the rate A3 states conditions only the modes above the Mosco scale: the conditioned cloud is Poisson at the grain in every frame and hyperuniform only at the constraint's scale (Computation 211). The one explicit placement the book has, classical scaling of the signed-tension metric, cannot test it, its Gram form being a negative multiple of the direction Gram matrix at leading order (Computation 209); the condition is supplied by the clause, not by a placement.

A4 (Isotropy / covariance forcing). The covariance matrix of the displacement $\rho_n(a)^\mu - x^\mu$ satisfies eq. (56):

$$\frac{1}{|D_n|}\sum_{a\in D_n}\big(\rho_n(a)^\mu - x^\mu\big)\big(\rho_n(a)^\nu - x^\nu\big) \tag{233}$$

$$\longrightarrow\ \frac{d_0^2}{4}\,g^{\mu\nu}(x). \tag{234}$$

Follows from A2 plus the absence of any preferred direction in V_7 surviving the configuration-space integral under μ. *Structural.*

A5 (Regularity of the Landau–Ginzburg quartic potential). The

Landau–Ginzburg free energy $F[\psi, \varepsilon]$ of (1) is quadratic at leading order around $\mathcal{V}$ and has bounded Hessian on bounded domains. *Structural from P3.*

A6 (Density of pulled-back test functions). The pullback

$$\Pi_n^\dagger C_c^\infty(M) \;\subset\; L^2(\mathcal{P}(D_n), \mu_n) \tag{235}$$

is dense in the form domain of E_n^B as $|D_n| \to \infty$. *Standard but technical*: follows from A2 plus the Stone–Weierstrass density of $C_c^\infty(M)$ in $L^2(M, dvol_g)$, combined with the explicit construction of $\Pi_n^\dagger$ via spherical-harmonic partial sums on S^3. A full proof requires explicit convergence-rate control of the truncation error.

A7 (Equicoercivity / Poincaré-type inequality). There exists a constant $C > 0$, independent of n, such that for every $u \in C_c^\infty(M)$,

$$\frac{1}{|D_n|}\, E_n^B(\Pi_n^\dagger u, \Pi_n^\dagger u) \;\geq\; C\, \|u\|^2_{H^1(M)}. \tag{236}$$

A theorem given A1, A2 and A4 together with A3, the last supplied by P3: the reduction lemma of §18.1.6 derives A7 with $C_K > 0$ from A1, A2, A4 together with A3. A3 is the rate hypothesis on the realised cloud (see the A3 entry above). Equicoercivity for the Boolean Dirichlet form is the discrete analogue of a Poincaré inequality on the substrate configuration space.

18.1.3. The Mosco-convergence theorem

18.1.3.1. Theorem (Mosco convergence, under A1–A7). *Under assumptions A1–A7 above, the rescaled Boolean Dirichlet forms* $\left(\frac{1}{|D_n|} E_n^B\right)$ *converge in the Mosco sense to* $\frac{d_0^2}{4}\, E^{\mathrm{Riem}}$ *on* $L^2(M, dvol_g)$*, where* d_0 *is the fixed coherence length (the A4 displacement covariance width, Chapter 15), unrelated to the vanishing site spacing* a_n *of A1:*

1. *(Liminf inequality).* For every sequence $u_n \in L^2(\mathcal{P}(D_n), \mu_n)$ with $u_n \to u$

in the limit identification implied by A2,

$$\liminf_{n\to\infty} \tfrac{1}{|D_n|}\, E_n^B(u_n, u_n) \;\geq\; \tfrac{d_0^2}{4}\, E^{\mathrm{Riem}}(u, u). \tag{237}$$

2. *(Recovery sequence).* For every $u \in H^1(M)$ there exists $u_n \in L^2(\mathcal{P}(D_n), \mu_n)$ with $u_n \to u$ and

$$\lim_{n\to\infty} \tfrac{1}{|D_n|}\, E_n^B(u_n, u_n) \;=\; \tfrac{d_0^2}{4}\, E^{\mathrm{Riem}}(u, u). \tag{238}$$

Throughout this chapter E^{Riem} and $\|\cdot\|_{L^2(M,g)}$ are normalised by the volume V_M, the normalisation A2 carries, so that $\frac{d_0^2}{4}E^{\mathrm{Riem}}$ here is the $\frac{d_0^2}{4V_M}\int_M$ of (57); Computation 177 measures the factor.

By Mosco's theorem [98], this implies strong resolvent convergence of the generators, $\big(\frac{1}{|D_n|}\, L_n^B + \lambda\big)^{-1} \Pi_n^\dagger \to \Pi\big(\frac{d_0^2}{4}\, \Delta_g + \lambda\big)^{-1}$ for all $\lambda > 0$, and the generator of the limiting form is the Laplace–Beltrami operator Δ_g, a Laplace-type operator with principal symbol $g^{\mu\nu}$.

18.1.3.2. Proof sketch of the Mosco theorem. *Liminf inequality.* Given $u_n \to u$, write $E_n^B(u_n, u_n) = \sum_{a\in D_n} \int (\partial_a^B u_n)^2\, d\mu_n$. By A2 and the matched scaling A1, the sum over $a \in D_n$ approximates the integral over M with respect to $dvol_g$. By A4, the covariance matrix forces the leading order of $(\partial_a^B u_n)^2$ pulled back to $x = \rho_n(a)$ to equal $(d_0^2/4) g^{\mu\nu}\partial_\mu u \partial_\nu u$, with d_0 the fixed displacement covariance width. Together, $\frac{1}{|D_n|} E_n^B(u_n, u_n) \to \frac{d_0^2}{4} E^{\mathrm{Riem}}(u, u)$. The *liminf* (rather than equality) version requires lower semicontinuity of $u_n \mapsto E_n^B(u_n, u_n)$ under A2-convergence, which follows from A7.

Recovery sequence. Given $u \in H^1(M)$, set $u_n = \Pi_n^\dagger u$. By A6, $u_n \to u$. By A2, A4 and the construction of $\Pi_n^\dagger$, $\frac{1}{|D_n|} E_n^B(u_n, u_n) \to \frac{d_0^2}{4} E^{\mathrm{Riem}}(u, u)$ with equality at the limit. A standard density argument extends from $C_c^\infty(M)$ to $H^1(M)$.

Strong resolvent convergence. Apply Mosco's theorem [98]; the generator of the limiting form is $\frac{d_0^2}{4}\Delta_g$ by direct identification.

The standing of A1 to A7 is as their entries above record: A1 is architectural

input; A2, A4 and A5 are structural; A3 is supplied by P3's maximum-entropy sublimation; A6 is technical but standard, a full proof wanting the spherical-harmonic-partial-sum projector with convergence-rate control; and A7 (equicoercivity) is a theorem given A1, A2, A4 and A3, via the reduction lemma of §18.1.6.

18.1.4. Numerical verification on finite models

Computation 4 tabulates the continuum target of the spin-structure component of the Mosco limit on $\mathbb{R} \times S^3$: the Friedrich–Bär Dirac spectrum $\pm(n+3/2)$ on S^3 and its Weyl law, the anchor that the matched scaling of Lemma 8 lands on (Computation 62); the finite-$|D|$ recovery is Lemma 8's. Computation 62 verifies the spectral-action invariance at finite $|D|$ as a consistency check on the matched-scaling convergence. No discrete finite-$|D|$ test currently corroborates A3 across $|D|$ for the emergence-determined cloud; such a test would be a natural concrete corroboration on the analytic side.

18.1.5. What the theorem establishes

The theorem above converts the heuristic statement "Mosco convergence yields $M = \mathbb{R} \times S^3$ as the Laplace-type limit" (used in §7.8 and downstream) into a precise theorem whose equidistribution condition A3 is supplied by P3's maximum-entropy sublimation. Under A1–A7, the strong-resolvent convergence is rigorous and the Laplace–Beltrami operator is identified as the generator of the limiting form, so the Mosco limit closes under A1–A7. It is not a theorem of P1–P3 alone: A1 is architectural input, and A6 still wants explicit convergence-rate control of the truncation error.

Equicoercivity A7 follows from A3 by the reduction lemma; the i.i.d. surrogate corroborates the rate (Computations 70, 83), and the corresponding discrete-geometry analytic quantification, whose continuum analogue is standard [113, 99], is a concrete deliverable rather than a gap. Consequently the downstream statements that depend on the Mosco limit (Einstein–Hilbert action from Seeley–DeWitt, the Connes spectral-triple structure of M, the Casimir prediction insofar as it relies on the kernel structure of $\Pi(L^B)$) inherit

that standing, conditional on A1–A7.

The structural content of PST does not depend on which level of analytic rigour one prefers, and the explicit assumption block makes the inputs, and their supply by P1–P3, transparent.

18.1.6. Sharpening A7: a uniform-rate inequality

The form of A7 stated above ("$\exists\, C > 0$ s.t. $(1/|D_n|)\, E_n^B(\Pi_n^\dagger u, \Pi_n^\dagger u) \geq C\, \|u\|_{H^1}^2$ for every $u \in C_c^\infty(M)$, independent of n") is the canonical Poincaré-type bound a Mosco-type proof would require. Under A1–A4 the proof-sketch argument above already gives the *pointwise* convergence

$$\frac{1}{|D_n|}\, E_n^B(\Pi_n^\dagger u, \Pi_n^\dagger u) \xrightarrow[n\to\infty]{} \frac{d_0^2}{4}\, \|\nabla u\|_{L^2(M,g)}^2, \qquad u \in C_c^\infty(M), \tag{239}$$

which together with the standard Poincaré inequality for compactly-supported smooth functions reduces A7 to a uniformity question. We state the reduction explicitly.

18.1.6.1. Lemma (A7 reduction). *Assume A1, A2 and A4 above, which give the pointwise convergence* (239)*. The equidistribution condition A3, supplied by P3's maximum-entropy sublimation, is precisely the statement that for each compact $K \subset M$ and each $R > 0$ there exist $N_0(K, R)$ and $\eta_{K,R} > 0$ with $\eta_{K,R} < d_0^2/4$ such that for all $n \geq N_0(K, R)$ and all $u \in C_c^\infty(K)$ with $\|u\|_{H^1(M)} \leq R$,*

$$\left| \frac{1}{|D_n|}\, E_n^B(\Pi_n^\dagger u, \Pi_n^\dagger u) - \frac{d_0^2}{4}\, \|\nabla u\|_{L^2(M,g)}^2 \right| \;\leq\; \eta_{K,R}\, \|u\|_{H^1(M)}^2. \tag{240}$$

Then A7 holds with K-dependent constant $C_K = \big(d_0^2/4 - \eta_{K,R}\big) \,/\, \big(C_P(K)^2 + 1\big)$, where $C_P(K)$ is the Poincaré constant of $K \subset (M, g)$ for test functions supported in K.

Proof. Standard Poincaré on $C_c^\infty(K)$, $\|u\|_{L^2(K)} \leq C_P(K)\, \|\nabla u\|_{L^2(K)}$ [113, 99],

gives $\|u\|_{H^1}^2 = \|u\|_{L^2}^2 + \|\nabla u\|_{L^2}^2 \le (C_P(K)^2+1)\,\|\nabla u\|_{L^2}^2$. Combining with (240):

$$\frac{1}{|D_n|}\,E_n^B(\Pi_n^\dagger u, \Pi_n^\dagger u) \;\ge\; \frac{d_0^2}{4}\,\|\nabla u\|_{L^2}^2 - \eta_{K,R}\,\|u\|_{H^1}^2 \;\ge\; \frac{d_0^2/4 - \eta_{K,R}}{C_P(K)^2+1}\,\|u\|_{H^1}^2, \tag{241}$$

which is A7 with constant $C_K = (d_0^2/4 - \eta_{K,R})/(C_P(K)^2+1) > 0$. □

18.1.7. The physical content of the bound

The equicoercivity constant $C_K = (d_0^2/4 - \eta_{K,R})/(C_P(K)^2+1)$ is not mere bookkeeping; it is the margin by which the emergent geometry stays non-degenerate. $C_K > 0$, equivalently the leading-coefficient inequality $\eta_{K,R} < d_0^2/4$, is exactly the condition that the rescaled Boolean Dirichlet form remains coercive in the Mosco limit, hence that the limit is a genuine continuum Dirichlet form rather than a collapsed one, the analytic statement that spacetime closes as a manifold. The rate condition is then not bookkeeping but a quantitative threshold: it measures how much equidistribution defect of the emergence map ρ_n, how much *unevenness* in the realised node measure against the Riemannian volume $\sqrt{-g}\,\mathrm{d}^4x$, the continuum geometry can absorb while still closing. As $\eta_{K,R} \uparrow d_0^2/4$ the constant $C_K \downarrow 0$ and the coercive lower bound is lost; past the threshold the construction no longer secures a non-degenerate limit. Read this way the rate condition is physical content: the equidistribution bound is the closure margin of emergent spacetime, the condition asks precisely where the T-determined cloud sits relative to it, and P3's maximum-entropy sublimation answers the question: the realised configuration is a typical one under μ, its cloud equidistributed at the rate A3 states, so the realised cloud sits inside the margin. Concretely, this residual unevenness is texture at the node (site) scale a_n: the permutation symmetry of the substrate measure evens it in measure, and the discreteness of the realised cloud smooths it away in the continuum limit, so the condition is precisely the rate at which that node-scale texture is erased.

The lemma converts the A7 problem from "find a uniform Poincaré-type spectral gap for the rescaled Boolean Laplacian" into "establish a uniform-

rate quantitative version of the pointwise-in-u convergence already known under A1–A4". Three points on the resulting reduced problem follow: its K-uniformity, its analytic content, and its numerical verification.

18.1.8. K-uniformity of the constant

The lemma delivers A7 with a constant C_K that depends on the support $K \subset M$ of the test function through the Poincaré constant $C_P(K)$ of K. Because $M = \mathbb{R} \times S^3$ is non-compact (a finite-volume compact factor S^3 times a non-compact $\mathbb{R}$), the Poincaré constant is bounded on any compact K but can grow without bound as K exhausts M in the temporal direction. For the downstream Mosco theorem only test functions of *compact* support are needed (recovery sequences are constructed by mollified cutoff arguments in $C_c^\infty(M)$ [98]), so the K-dependence is not an obstruction. A genuinely K-uniform constant would require a stronger statement than is needed for the downstream applications.

18.1.9. The analytic content of the rate hypothesis

The pointwise convergence (239) relies on the law-of-large-numbers limit in A2 and the covariance identity A4. Both have explicit finite-n correction rates for the i.i.d. surrogate (Berry–Esseen rate $|D_n|^{-1/2}$ [101]). For the emergence-determined ρ_n the *fluctuation* of the empirical aggregate A_n (equation (242)) carries the same $|D_n|^{-1/2}$ rate by McDiarmid concentration on the i.i.d. bits under a single-bit-flip locality hypothesis (§18.2, no statistical independence of positions assumed; Computation 125); the *bias* (mean) rate, with the A4 covariance-scale convergence, is the content that P3's maximum-entropy sublimation supplies (corroborated on the i.i.d. surrogate).

The reduced problem (240) is therefore quantitative: it asks for a Berry–Esseen-type rate of order $|D_n|^{-\gamma}$ for some $\gamma > 0$, uniform on the H^1-ball of $C_c^\infty(K)$. For the i.i.d. surrogate this is a discrete-analytic statement of the same form as the finite-element rate estimates for elliptic PDEs [102], closed at proof-detail level by the five-step Berry–Esseen argument below. For the emergence-determined ρ_n it is sharper than an exponent: the lemma needs

the *leading-coefficient* inequality $\eta_{K,R} < d_0^2/4$ (not merely $\eta_{K,R} \to 0$), so a universality result giving only the rate γ does not suffice.

A covariance-decay (spectral-gap) mixing estimate is, however, the wrong instrument for this residual. A spectral gap is a relaxation *rate*, and bounds only the statistical fluctuation of A_n about its mean; the stationary mean is the kernel eigenvector, decoupled from the gap. No mixing exponent therefore reaches the leading coefficient. Computation 129 shows this directly: the gap is invariant under a full sweep of the stationary mean, while the systematic coefficient is flat in $|D_n|$ as the fluctuation decays at $|D_n|^{-1/2}$. The binding content is therefore the *systematic* mean deviation: a deterministic equidistribution defect of the emergence-determined ρ_n in the invariant sector $\ker L$. What is needed is the leading-coefficient (equidistribution) inequality $\eta_{K,R} < d_0^2/4$ itself, not a quantitative ergodic theorem. The spectral-gap and mixing route is therefore foreclosed for the coefficient, as pinning T and local dependence are, and it is exactly this invariant-sector coefficient that P3's maximum-entropy sublimation supplies.

This equidistribution defect is concretely the uniform-weight nodal quadrature error of the emergence cloud against the continuum measure. Writing $c_u = |D_n|^{-1} \sum_a |\nabla u(\rho_n(a))|^2 \big/ \big(|K|^{-1} \int_K |\nabla u|^2\big)$, one has $c_u \to 1$ iff the realised points cell-volume-equidistribute: the deterministic sibling of the A2 weak (H^1-tested) discrepancy. It is *not* the Delaunay finite-element interpolation error, which mesh shape-regularity would control; shape-regularity is here neither necessary nor sufficient. The matched-scaling assumption A1 does not supply it: A1 fixes only the average density $|D_n|/V_M \sim a_n^{-4} \to \infty$, not the minimum separation or the local equidistribution, and a matched-count cloud carrying a void or density excess retains a nonzero c_u (Computation 130). The strong separation channel, by contrast, does not bind the weak discrepancy. The directed metric (30) manufactures vanishing-separation pairs on the null cone, but that locus is codimension ≥ 1 on pair-space, and leaves the single-node empirical measure, hence c_u, unchanged (Computation 131). It therefore does not obstruct A7.

The bias half therefore rests on a single named deterministic gauge-invariant lemma, the codimension-zero cell-volume equidistribution (uniform bulk fibre-occupation of the node measure over the 4-volume) of the T-determined cloud, which matched scaling does not force and which is precisely what P3's maximum-entropy sublimation supplies. Five constructions are each provably blind to it: the globally-rigid quotient embedding, the displacement-covariance operator, the Lorentzian (30), the Mosco recovery sequence, and the tension functional itself. A positive-volume void leaves every pairwise tension value, the magnitude ordering of $|T(C)|$, and the global bias $\hat{\tau}$ invariant (Computation 132), so it is invisible to all three channels through which T reaches physics. The theory's one genuinely variational term, the Landau–Ginzburg gradient energy, is minimised over the modal order parameter on the configuration lattice, never over spatial node density.

Permutation symmetry of the substrate measure forces equidistribution *in measure*, which forbids a systematic asymmetric bias but supplies neither the positive-rate weak (H^1-tested) discrepancy nor the leading-coefficient inequality $\eta_{K,R} < d_0^2/4$, the precise, fully-located content beyond symmetry. This is a gauge-invariant analytic *rate*, not a new postulate: the positive-rate weak (H^1-tested) discrepancy of the T-determined ρ_n against the Riemannian volume $\sqrt{-g}\,\mathrm{d}^4x$ is supplied by P3's maximum-entropy sublimation and corroborated on the i.i.d. surrogate.

A diagnosis in terms of an irreducible coordinate bulk-density freedom would be a category error: the node-quadrature c_u above carries *flat* weights and is therefore chart-dependent, so a metric-preserving (diffeomorphic) void biases it only by dropping the chart's own g^{-1} and $\sqrt{g}$ factors, whereas the metric-covariant defect is identically void-blind (Computations 133 and 134); a void and its co-transforming metric form a single gauge orbit and carry no independent freedom. A7 therefore needs no bulk-density axiom and no new postulate beyond P1–P3. It is the intrinsic, gauge-invariant equicoercivity of the Boolean Dirichlet form E_n^B, of which A_n is the exact pullback value (Closure (a)), resting on one diffeomorphism-invariant equidistribution rate.

That rate is the standard quantitative-realisation input of any discrete-to-continuum limit, supplied for the emergence-determined ρ_n by P3's maximum-entropy sublimation and corroborated on the i.i.d. surrogate. A genuine fixed-metric density modulation, as opposed to a coordinate void, would bias the metric-covariant defect (Computation 134); the μ-typical configuration carries none.

Two honest provisos attach to the rates just quoted. The displacement is modelled as i.i.d. Gaussian of covariance $(d_0^2/4)g$, the A4 covariance forcing, and the $|D_n|^{-1/2}$ rates use the i.i.d. structure of the realisation. Permutation invariance of μ delivers equidistribution *in measure* but not statistical independence, so the quantitative forms of A2 (a positive-rate weak (H^1-tested) discrepancy of ρ_n) and A4 (a positive-rate covariance isotropy of ρ_n) are not consequences of symmetry alone. They are realisation properties of the μ-typical configuration, exactly what P3's maximum-entropy sublimation asserts of the instantiated substrate, and the i.i.d. surrogate is the vehicle on which they are corroborated (Computations 70, 83).

18.1.10. Numerical verification of the rate

Computation 70 verifies (240) numerically on an explicit flat-bulk toy model. Three compactly-supported smooth test functions of bounded H^1-norm (Gaussian, C_c^∞ bump, $(1-r^2)^2\cos(\pi r^2)$) are tested at $|D_n| \in \{10^4, \ldots, 3\times 10^6\}$. Two findings:

1. The empirical std-error of $A_n(u)$ across independent realisations scales as $|D_n|^{-1/2}$ (empirical exponent -0.467 for the Gaussian, within the Berry–Esseen prediction -0.5). The statistical part of $|A_n - A_\infty|$ is exactly the central limit theorem (CLT) rate predicted above.

2. The systematic part of $|A_n - A_\infty|$ scales as $O(d_0^4)$ in the fixed coherence length d_0, the next-order Taylor correction beyond the leading covariance forcing A4. It is a fixed finite-d_0 bias, small in the physical regime $\mathrm{radius}(K) \gg d_0$ (relative correction $O((d_0/\mathrm{radius}(K))^2)$), and is bounded by $\|u\|_{H^1}^2$ with an explicit constant involving $\|\mathrm{Hess}(u)\|$.

Together these supply, on the Gaussian-bump-trigonometric class, the systematic error $O(d_0^4 \, \|\mathrm{Hess}\, u\|^2)$ and statistical error $O(d_0^2 \, |D_n|^{-1/2})$, i.e. a *relative* rate $\eta_{K,R}/(d_0^2/4) = O\big((d_0/\mathrm{radius}(K))^2 + |D_n|^{-1/2}\big)$, small in the physical regime $\mathrm{radius}(K) \gg d_0$ and as $|D_n| \to \infty$, confirming the analysis above quantitatively. Under the rate hypothesis with a relative rate below unity the Lemma above closes A7, with a K-dependent constant, which is all the recovery-sequence step requires.

18.2. The K-uniform rate: proof and three closures

This section proves, for the i.i.d. surrogate, the uniform-rate hypothesis of §18.1, by the five-step Berry–Esseen argument and the three closures (a)–(c); from it, assumption A7 of the Mosco theorem follows by the reduction lemma proved there. We give the analytic argument for the rate hypothesis in the natural physical regime (test-function support large relative to the substrate coherence d_0), using standard Brenner–Scott finite-element machinery for the discrete-to-continuum estimate. The argument is presented for the flat-bulk case $M = \mathbb{R}^3$ where the explicit constants are transparent; extension to $M = \mathbb{R} \times S^3$ follows by replacing the Sobolev embedding constants with the Riemannian analogues, which differ by bounded geometric factors.

18.2.1. Setup: the domain K and the empirical aggregate

Let $K \subset \mathbb{R}^3$ be a bounded open set with Lipschitz boundary, $u \in C_c^\infty(K)$, and $\Pi_n^\dagger u \in L^2(\mathcal{P}(D_n), \mu_n)$ the pullback via the quasi-interpolation construction implicit in the proof-sketch of Computation 70. The key empirical aggregate is

$$A_n(u) \;=\; \frac{1}{|D_n|} \sum_{a \in D_n} \big(u(\rho_n(a) + \xi_a) - u(\rho_n(a))\big)^2, \tag{242}$$

with ξ_a i.i.d. Gaussian on $\mathbb{R}^3$ with covariance $(d_0^2/4)\, I_3$.

18.2.2. Step 1: Taylor expansion

For $u \in C^3(K)$ and $|\xi_a| \ll \operatorname{dist}(\rho_n(a), \partial K)$,

$$\begin{aligned} u(\rho_n(a)+\xi_a) - u(\rho_n(a)) \;&=\; \xi_a \cdot \nabla u + \tfrac{1}{2}\xi_a^T \mathrm{Hess}(u)\xi_a \\ &\quad + O(|\xi_a|^3 \cdot \|D^3 u\|_\infty). \end{aligned} \tag{243}$$

Squaring,

$$\begin{aligned} \big(u(\rho_n(a)+\xi_a) - u(\rho_n(a))\big)^2 = (\xi_a \cdot \nabla u)^2 + \xi_a \cdot \nabla u \cdot \xi_a^T \mathrm{Hess}(u)\xi_a \\ + O\big(|\xi_a|^4 \cdot \|D^3 u\|_\infty^2 \\ + |\xi_a|^4 \cdot \|\mathrm{Hess}(u)\|^2\big). \end{aligned} \tag{244}$$

18.2.3. Step 2: Expectation under ξ

Using $\mathbb{E}[\xi_a \xi_a^T] = (d_0^2/4) I_3$ (a positive Euclidean covariance, not the indefinite metric; Computation 124), isotropy $\mathbb{E}[(\xi \cdot v)^2] = (d_0^2/4)|v|^2$ for any v, and the fact that odd-power moments of an isotropic Gaussian vanish:

$$\begin{aligned} \mathbb{E}\big[(u(\rho_n(a)+\xi_a) - u(\rho_n(a)))^2\big] \;&=\; \tfrac{d_0^2}{4}\,|\nabla u(\rho_n(a))|^2 \\ &\quad + O\big(d_0^4 \cdot \|\mathrm{Hess}(u)\|^2 \\ &\quad + d_0^6 \cdot \|D^3 u\|_\infty^2\big). \end{aligned} \tag{245}$$

18.2.4. Step 3: Empirical average over the realisation map

Separate the *bias* (mean) from the *fluctuation*, $A_n(u) = \mathbb{E}[A_n(u)] + (A_n(u) - \mathbb{E}[A_n(u)])$, the expectation taken over the independent inputs (the i.i.d. Bernoulli bits the realisation ρ_n is built from, and the stipulated i.i.d. displacements ξ_a).

Bias. Averaging the Step-2 expectation over $a \in D_n$,

$$\mathbb{E}[A_n(u)] \;=\; \tfrac{d_0^2}{4} \cdot \tfrac{1}{|D_n|} \sum_{a \in D_n} |\nabla u(\rho_n(a))|^2 + O(d_0^4 \cdot \text{Hess-sup}_K(u)^2), \tag{246}$$

and $\frac{1}{|D_n|}\sum_a |\nabla u(\rho_n(a))|^2 \to \frac{1}{|K|}\int_K |\nabla u|^2$ by the equidistribution A2. The *rate* of this limit is the weak (H^1-tested) discrepancy of the emergence-determined cloud $\{\rho_n(a)\}$; it is not supplied by the present (surrogate) argument but by P3's maximum-entropy sublimation, as the equidistribution condition A3 (§18.1.6).

Fluctuation. $A_n(u) = \Psi_n(x)$ is a deterministic function of the independent inputs x above. Provided the realisation map has the *bounded-differences* property (flipping one bit relocates $O(1)$ points by $O(d_0)$, hence changes A_n by $c_a = O(\|u\|^2_{C^1}\, d_0^2/|D_n|)$ per coordinate), McDiarmid's inequality [100] gives

$$\mathbb{P}\big(|A_n - \mathbb{E}A_n| > t\big) \;\le\; 2\exp\Big(-\frac{2t^2}{\sum_a c_a^2}\Big), \qquad \sum_a c_a^2 = O\big(d_0^4/|D_n|\big), \quad (247)$$

so the fluctuation concentrates at scale $O(d_0^2|D_n|^{-1/2})$. This is the honest replacement for an i.i.d. central-limit sum: it does *not* require the positions $\{\rho_n(a)\}$ to be statistically independent (which permutation invariance of μ does not supply), only the weaker single-bit-flip locality of the realisation map, a structural property of the emergence construction. Computation 125 confirms this single-bit-flip locality and the $O(|D_n|^{-1/2})$ fluctuation rate on an explicit deterministic-functional-of-i.i.d.-bits toy, with the bias channel correspondingly global.

18.2.5. Step 4: K-uniform bounds

The McDiarmid increment $c_a = O(\|u\|^2_{C^1} d_0^2/|D_n|)$ and, for the per-term scale, the variance $\mathbb{V}[X_a]$ and third absolute moment $\mathbb{E}[|X_a|^3]$ (the latter retained only for the i.i.d. surrogate central-limit rate comparison; the K-uniform rate of Step 5 rests on the McDiarmid concentration scale of Step 3, not on the third moment) are bounded by quantities depending on $\|u\|_{C^3(K)}$ and on d_0:

$$\mathbb{V}[X_a] \;\le\; C \cdot d_0^4 \cdot \|u\|^2_{C^1(K)} \cdot \|u\|^2_{C^2(K)}, \quad (248)$$

$$\mathbb{E}[|X_a|^3] \;\le\; C' \cdot d_0^6 \cdot \|u\|^6_{C^1(K)} + C'' \cdot d_0^{12} \cdot \|u\|^6_{C^3(K)}. \quad (249)$$

For u bounded in $H^1(K)$ with bounded $\|\mathrm{Hess}(u)\|$ (which is automatic for smooth test functions on bounded K), these are uniformly bounded across the family $\mathcal{F}(K,R) = \{u \in C_c^\infty(K) : \|u\|_{H^1(M)} \le R\}$.

18.2.6. Step 5: K-uniform rate

Combining Steps 3 and 4, for $u \in \mathcal{F}(K,R)$ and n large enough,

$$\left|A_n(u) - \tfrac{d_0^2}{4}\|\nabla u\|_{L^2}^2\right| \;\le\; C(K,R)\,\big(d_0^4\,\|\mathrm{Hess}\,u\|_{L^2(K)}^2 + d_0^2\,n^{-1/2}\,\|u\|_{H^1}^2\big), \tag{250}$$

with $C(K,R)$ depending on K only through the Poincaré constant $C_P(K)$ (standard for compact K with Lipschitz boundary [102], Ch. 4) and the Sobolev-embedding constants of K (bounded for compact K). Dividing by the leading term $\frac{d_0^2}{4}\|\nabla u\|_{L^2}^2$ gives the *relative* rate $\eta_{K,R}/(d_0^2/4) = O\big((d_0/\mathrm{radius}(K))^2 + n^{-1/2}\big)$, small in the physical regime $\mathrm{radius}(K) \gg d_0$ (fixed coherence length) and as the site count $|D_n| = n \to \infty$ at site spacing $a_n \to 0$; the bound is uniform over the Hessian-bounded class $\{u \in C_c^\infty(K) : \|u\|_{H^1} \le R,\ \|\mathrm{Hess}\,u\|_\infty \le R'\}$, which contains the mollified cutoff functions the Mosco recovery sequences use, so the restriction from the full H^1-ball is not felt downstream. This is the K-uniform rate of §18.1.6.

18.2.7. Caveats and what remains

1. *Thin-support regime.* For u with $\mathrm{supp}(u)$ of radius comparable to d_0, the higher-order Taylor terms in Step 1 dominate and the rate must be supplemented by explicit d_0-dependent corrections. This is the regime of Computation 83's failing families and is not closed by the present proof.

2. *Riemannian extension.* The full case $M = \mathbb{R} \times S^3$ requires replacing the flat-bulk Sobolev embedding constants by the Riemannian analogues on S^3. These differ by bounded geometric factors that do not affect the rate but do affect the explicit $C(K,R)$. The argument extends *verbatim* with the substitution; we have not written out the explicit constants here.

3. *Explicit* $\Pi_n^\dagger$. The proof above works at the level of the empirical aggregate $A_n(u)$; an explicit $\Pi_n^\dagger$ construction that realises this aggregate as $(1/|D|)\, E_n^B(\Pi_n^\dagger u, \Pi_n^\dagger u)$ is presupposed. Standard quasi-interpolation constructions on the substrate would do this; we have not constructed one explicitly here.

18.2.8. Closure (a): Explicit $\Pi_n^\dagger$ construction

For $u \in C_c^\infty(K)$ with $K \subset M$ compact, the natural quasi-interpolation operator $\Pi_n^\dagger : C_c^\infty(K) \to L^2(\mathcal{P}(D_n), \mu_n)$ is built from the substrate's Delaunay triangulation. Given the realisation $\rho_n : D_n \to M$ (asymptotically equidistributed by A2), the Delaunay triangulation $\mathcal{T}_n$ of $\{\rho_n(a)\}_{a\in D_n} \cap K$ is the standard simplicial cover with mesh size $h_n \asymp d_0$ (by the matched scaling A1). Let $\{\phi_a\}_{a\in D_n}$ be the corresponding hat-function basis on $\mathcal{T}_n$, the unique piecewise linear function with $\phi_a(\rho_n(b)) = \delta_{ab}$, and define

$$\begin{aligned}(\Pi_n^\dagger u)(C) \;:=\; & \frac{1}{\sqrt{|D_n|}} \sum_{a\in C} \phi_a\big(\rho_n(a)\big)\cdot u\big(\rho_n(a)\big) \\ & - \frac{1}{\sqrt{|D_n|}} \sum_{a\in D_n\setminus C} \phi_a\big(\rho_n(a)\big)\cdot u\big(\rho_n(a)\big),\end{aligned} \tag{251}$$

which is the Lagrange quasi-interpolation in the Walsh basis on the hypercube $\mathcal{P}(D_n)$. By construction $\phi_a(\rho_n(a)) = 1$, so (251) simplifies to

$$(\Pi_n^\dagger u)(C) \;=\; \tfrac{1}{\sqrt{|D_n|}} \sum_{a\in D_n} (-1)^{\mathbf{1}[a\notin C]}\, u(\rho_n(a)), \tag{252}$$

i.e. the Walsh–Fourier sum of $u(\rho_n(a))$ along single-site weights. The Boolean derivative $\partial_a^B(\Pi_n^\dagger u)(C) = (-2/\sqrt{|D_n|})\, u(\rho_n(a))$, and the Boolean Dirichlet form

$$\frac{1}{|D_n|}\, E_n^B(\Pi_n^\dagger u, \Pi_n^\dagger u) \;=\; \frac{4}{|D_n|^2} \sum_{a\in D_n} u(\rho_n(a))^2. \tag{253}$$

Up to the constant factor 4, this is the empirical aggregate $A_n(u)$ of the proof sketch above with $\xi_a = 0$ (no displacement); it captures the L^2 norm but not the H^1 norm. The H^1-encoding $\Pi_n^\dagger$ requires the finite-difference version:

replace $u(\rho_n(a))$ by $u(\rho_n(a)+\xi_a)-u(\rho_n(a))$ where ξ_a is the displacement to a Delaunay-neighbour of a. Two normalisations must be fixed for this to be an identity rather than an assertion. (i) The Walsh quasi-interpolant carries a factor $\frac{1}{2}\sqrt{|D_n|}$ relative to eq. (251), chosen to cancel the spurious $4/|D_n|$ of the zero-displacement computation above, so that $(1/|D_n|)E_n^B(\Pi_n^\dagger u,\Pi_n^\dagger u)=A_n(u)$ holds. (ii) The displacement ξ_a must carry covariance $(d_0^2/4)g$ (the A4 forcing), with d_0 the fixed coherence length; a literal Delaunay-edge vector, whose length is the mean edge length $\ell_n \asymp a_n \to 0$ (the vanishing site spacing of A1), must be rescaled by the fixed factor d_0/ℓ_n so that its covariance matches the fixed coherence scale $(d_0^2/4)g$. With these, the empirical aggregate $A_n(u)$ of equation (242) is realised and the proof sketch delivers the K-uniform rate.

18.2.9. Closure (b): Thin-support d_0-corrections

For test functions u with support of radius $\sim d_0$, the Taylor expansion of equation (243) keeps only the second-order term $\xi^T\mathrm{Hess}(u)\xi/2$ with relative error $O(d_0\|D^3u\|_\infty/\|D^2u\|_\infty)=O(1)$ in the thin-support regime. The leading d_0-correction comes from the third-order Taylor term and is bounded by

$$\begin{aligned}|A_n(u)-\tfrac{d_0^2}{4}\,\|\nabla u\|_{L^2}^2|\;\le\; C\Big(&d_0^2\cdot\text{Hess-sup}_K(u)\\&+d_0^4\cdot D^3\text{-sup}_K(u)\Big)\,\|u\|_{H^1}^2,\end{aligned}\tag{254}$$

where the constant C depends on the Poincaré constant of K but is d_0-independent (Brenner–Scott [102], Ch. 4). In the asymptotic regime $\mathrm{radius}(K)\gg d_0$ the second-order Hessian term dominates and the rate is $O(d_0^2)$; in the thin-support regime both terms contribute with comparable weight. This recovers Computation 83's empirical finding that the failing families precisely match the thin-support condition.

18.2.10. Closure (c): Riemannian extension on $M=\mathbb{R}\times S^3$

The proof sketch was carried out for $M=\mathbb{R}^3$ flat-bulk. Extension to $M=\mathbb{R}\times S^3$ proceeds by replacing each flat-Euclidean ingredient by its Riemannian analogue: the gradient ∇u in equation (250) is now the Riemannian gradient

$\nabla_g u$ with g the product metric $dt^2 + g_{S^3}$, where g_{S^3} is the round metric on S^3. The Sobolev embedding constants for $H^1(K) \hookrightarrow L^q(K)$ on compact $K \subset M$ are bounded by their flat-Euclidean counterparts up to a multiplicative factor of order 1 (Friedrich [113], Bär [114], Taylor [99]). The displacement ξ_a is replaced by a normal vector in $T_{\rho_n(a)}M$ with covariance $(d_0^2/4)\, g^{\mu\nu}(\rho_n(a))$. The Taylor expansion acquires Christoffel-symbol corrections at third order and beyond, all bounded uniformly on compact K. None of these substitutions affects the Berry–Esseen rate at leading order; they affect the explicit constant $C(K, R)$ by bounded geometric factors. The K-uniform Berry–Esseen rate therefore holds, for the i.i.d. surrogate, on $M = \mathbb{R} \times S^3$ with constants depending on K only through the Riemannian Poincaré constant of K and bounded curvature quantities of g on K.

The explicit $\Pi_n^\dagger$ is the Walsh-finite-difference version of the Delaunay Lagrange interpolation; the thin-support corrections are governed by equation (254); the Riemannian extension proceeds by standard substitution. Together with the proof sketch Steps 1–5 above, Closures (a)–(c) prove the rate hypothesis for the i.i.d. surrogate, corroborating the rate condition A3 that P3 supplies, and hence A7 via the reduction lemma of §18.1.6.

18.2.11. Explicit constants: Berry–Esseen and finite elements

The explicit-constant content referenced in the above closures is standard finite-element analysis. The classical Berry–Esseen bound Δ_{BE}, quoted here only for the i.i.d. surrogate central-limit comparison (the fluctuation channel of the proof is carried by the McDiarmid concentration of Step 3, not by an i.i.d. central-limit sum), is the optimal classical constant $\Delta_{\mathrm{BE}} = c_{\mathrm{BE}}\, \sigma^{-3}\, \rho_3/\sqrt{n}$ with $c_{\mathrm{BE}} = 0.4748$ (Shevtsova [104], sharpening Esseen [103]), where σ^2 is the displacement variance and ρ_3 is the third absolute moment of the per-site displacement. For the Sobolev embedding $H^1(K) \hookrightarrow L^2(K)$ on bounded Lipschitz $K \subset M = \mathbb{R} \times S^3$, the embedding constant is $C_S(K) = O\big((|K|)^{1/2}(1 + L_K)\big)$ with L_K the Lipschitz constant of ∂K; the Poincaré inequality gives $\|u - \bar{u}\|_{L^2(K)} \le C_P(K)\, \|\nabla u\|_{L^2(K)}$ with $C_P(K) \le \mathrm{diam}(K)/\pi$

for convex K (Brenner–Scott [102], Ch. 4). On the Riemannian $M = \mathbb{R} \times S^3$, the Bochner–Lichnerowicz formula [152, 153] gives a multiplicative curvature factor $(1 + C_g \cdot \kappa_{\max}(K))$ in the Sobolev embedding constant, where $\kappa_{\max}(K)$ is the bound on the Ricci curvature of g on K (Friedrich [113], Bär [114]) and C_g is a metric-dependent geometric constant of order 1. Together these give a fully explicit quantitative form of the K-uniform rate.

18.3. Dirac-operator convergence: a single 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

$$\big(C^\infty(\mathbb{R} \times S^3),\ L^2(\mathcal{S}),\ D_{\mathbb{R}\times S^3},\ J,\ \gamma\big) \tag{255}$$

in an appropriate substrate-to-continuum metric. The closure target for PST is the relaxed norm-equivalence criterion of equation (257) below (§18.3.5), in the distance framework of [116], strictly weaker than Latrémolière's spectral propinquity [117, 118] (operator-level convergence of Dirac data). The relaxation is structurally forced by the foundational object's fixed internal Hilbert space $H_F = \mathbb{C}^8$ (derived in §13.7). Computation 9 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.

18.3.1. The three ingredients of the Dirac estimate

1. *Mosco / strong-resolvent convergence of the Boolean Dirichlet forms to the S^3 Laplace–Beltrami form.* Established in §7.8; 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 [172] $w_2(M)$ 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). Friedrich's eigenvalue formula [113, 114] gives the S^3 Dirac spectrum $\pm(n+\frac{3}{2})$ with multiplicity $2(n+1)(n+2)$.

3. *Boolean–canonical-anticommutation-relations (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 PST closure target is the relaxed norm-equivalence form below (equation (257); §18.3.5).

18.3.2. Ingredient 3, what is exact at finite substrate size: the Clifford algebra, the multiplicities and the bridge map

At every finite substrate size D the Boolean–CAR generators χ_a, χ'_a ($a = 1, \ldots, D$) satisfy the $\mathrm{Cl}(0, 2D)$ anticommutation $\{\chi_a, \chi_b\} = 2\delta_{ab} = \{\chi'_a, \chi'_b\}$, $\{\chi_a, \chi'_b\} = 0$ *exactly* (Computation 9, section 2 of the script, machine-precision check up to $D = 6$); the Walsh-mode [165] binomial multiplicities $\binom{D}{k}$ flow under the Landau–Ginzburg (LG) induced reweighting (Chapters 4 and 7) to the $(k+1)^2$ spherical-harmonic multiplicities at every k in the thermodynamic limit (Computation 9, section 3 of the script, 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, section 7 of the script); 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, section 8 of the script).

18.3.3. Ingredient 3, what remains: the commutator residual as a Sobolev-tail estimate

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

$C^\infty(S^3)$:

$$\left\| [\chi_a^{(D)}, f_D] - b_D([\gamma_a, f]) \right\|_{\mathrm{op}} \;\le\; C\, D^{-s/2}\, \|f\|_{H^{s+1}(S^3)}, \qquad f \in H^{s+1}(S^3), \tag{256}$$

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, section 7 of the script, and γ_a is the continuum Clifford action. Computation 9, section 11 of the script 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 [165]) are standard Sobolev / Clifford-derivation manipulations on compact spin manifolds; the numerical check in Computation 9, section 13 of the script confirms the tail-rate scaling. Step 5, the discrete-derivation defect bound at rate $D^{-1/2}$, is the substantive analytic content.

18.3.4. Step 5, the discrete-derivation defect: why no citation closes it, and what does

A direct inspection of Latrémolière 2018 [117, 118] (Computation 9, section 15 of the script) confirms that step 5 is *not* closed by routine citation under the spectral-propinquity framing: spectral-propinquity convergence is established case-by-case for each finite-dimensional approximation scheme. The two sphere-approximation precedents in the literature (Rieffel's matrix-algebra $\to S^2$ convergence (Berezin/Kähler scheme, in the quantum Gromov–Hausdorff (QGH) distance, [115]) and Bhattacharyya–Singla's Toeplitz-algebra [144] $\to$ S^{2k+1} convergence (Bergman scheme, also in the quantum Gromov–Hausdorff distance, [119])) both use complex-projective / coherent-state machinery that does not directly port to the Boolean cube's real Dirichlet-form structure. The cutoff analysis of Diaconis and Shahshahani [120], for the Bernoulli–Laplace model and, by the same method, the random walk on the hypercube, gives the measure-level rate: stationarity is reached after of order $D \log D$

steps, so the per-step convergence rate is of order $(\log D)/D$; lifting this to operator-level convergence in the spectral-propinquity framework is not the route PST commits to. The closure target (§18.3.5) is the relaxed norm-equivalence criterion of equation (257), for which the Fréchet smooth Lip-norm machinery from L_round closure (Computations 20–23) transfers; the empirical $\gamma_D = 0.007$ at $D = 8$, $N = 20$ (an offline run; the published Computation 36 runs $N = 10$ for the browser) at the Walsh-weight cutoff (Computation 36) is the finite-truncation verification, with the analytical closure achieved via the refined symmetric-monomial bridge (Computations 39–42).

The Connes real spectral-triple identification has six checkable pieces. Five are verified at the structural level with no obstruction.

- Mosco resolvent convergence of the Boolean Laplacian (§7.8).
- Unique spin structure (Computation 4).
- Exact discrete Clifford algebra at every D (Computation 9, section 2 of the script).
- Exact multiplicity flow (Computation 9, section 3 of the script).
- Exact low-mode commutator gap (Computation 9, section 8 of the script).

The sixth piece, the Sobolev-tail bound of equation (256), is the residual. The residual splits cleanly into two sub-statements, an observation we record explicitly (see [116]). The first is a quantum Gromov–Hausdorff section round-trip lemma analogous to Bhattacharyya–Singla 2022 Lemma 3.2. The second is the lift of that round-trip from quantum Gromov–Hausdorff to spectral propinquity, the genuinely Dirac-aware statement of equation (256).

18.3.4.1. Section round-trip in Gromov–Hausdorff distance. The first sub-statement, the section round-trip, is the L_round side, and it closes rigorously. With $\pi_D : A_D \to C(S^3)$ the Walsh-coefficient truncation at weight $|S| \leq k_D := \lfloor \sqrt{D} \rfloor$ and σ_D the trivial section, the Fréchet smooth Lip-norm L_F of equation (259) satisfies the round-trip bound (258) at the rate $\gamma_D = O(D^{1/4} e^{-2\sqrt{D}})$, by the closed-form block decomposition (260) and

the exact off-diagonal norm $\|P_+ T_{\text{tail}} P_-\|_{\text{op}} = 2^{D-1} - D$ of equation (261) (Computations 20, 22, 23, 50–52, 54, 55); polynomial Walsh-weight Lip-norms fail because the constant-coefficient tail has operator norm 2^D while any polynomial weight grows only as D^α. Combined with the bridge construction of [116], this gives convergence $(\text{Lip}(A_D), L_F) \to (C(S^3), L)$ in the quantum Gromov–Hausdorff distance at exponential rate. The derivation in full, with the lemma chain and every singular-value computation, is §18.4.

18.3.5. Dirac-aware lift to substrate-to-continuum metric

With the section round-trip closed (rigorously, via the block decomposition), 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 ([117, 118]) at the level of equation (256), requiring full operator equality of bridge commutators. The actual closure target for PST (below) is the relaxed norm-equivalence criterion equation (257) in the [116] 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 route's first probes, the explicit substrate–Toeplitz bridge at $D = 4$, its holomorphic restriction, the optimal-α diagnostic and the first spinor extension, are Computations 28–31, kept in Appendix B as the record of how the non-tensor coupling was arrived at.

The discrete-side audit (Computation 21, Computation 24) confirms the algebraic structure in two ways. 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 $l = |S|$ ([113]; tabulated and cross-checked against Friedrich's eigenspinor decomposition in Computation 25). And $[\chi_a, \chi_S]$ is Hilbert–Schmidt [173] 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 is 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 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 implements the full BS 2022 Toeplitz framework on $H^2_\alpha(B^2)$ explicitly. The Toeplitz commutator $[T_{\bar{z}_1^k}, T_{z_1^k}]$ 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 $R(k, \alpha) = \|\text{comm}\|/\|T_{z^k}\|$ is therefore α-dependent, providing precisely the non-proportional α-scaling needed. The BS 2022 closure mechanism is thus explicitly verified on $H^2_\alpha(B^2)$: $\alpha = \alpha(D)$ supplies the internal degree of freedom needed to satisfy both fidelity and gap rate.

With the BS 2022 Toeplitz framework on $H^2_\alpha(B^2)$ explicitly verified (Computation 27), the explicit bridge identifying the PST substrate $\text{Cl}(0, 2D)$ algebra with the Toeplitz algebra is the non-tensor coupling of Computation 32: the round-S^3 Dirac $D = \sum_a \sigma_a \otimes J_a$ on $H^2_\alpha(B^2) \otimes \mathbb{C}^2$, where J_a are the SU(2) generators acting on the Bergman monomial basis as raising/lowering operators. The framework is operational: the Hermitian (in Hermite's sense [184]) Dirac operator is verified, the su(2) algebra holds to machine precision, and the eigenvalue range approximates Friedrich's $\pm(n + 3/2)$ S^3 spectrum. 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 and 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 and 0.87 at $k = 1, 2, 3$. Scaling to $D = 6$, $N = 12$ (an offline run; the published Computation 34 runs $N = 8$ for the browser) 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*:

$$\begin{aligned} \big|\, \|[D_\alpha, \text{bridge}(\chi_S)]\|_{\text{op}} & \\ & - \|[D_{\text{sub}}, \chi_S]\|_{\text{op}} \,\big| \\ \leq\ & \gamma_D \cdot \|\chi_S\|_{\text{op}}, \\ & \gamma_D \to 0 \text{ at exponential rate,} \end{aligned} \tag{257}$$

on the Walsh-weight-bounded subalgebra $|S| \leq k_D = \lfloor \sqrt{D} \rfloor$. The choice of this relaxed target 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 $H_F = \mathbb{C}^8$ (§13.7). 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 $\mathbb{C}^8$.

18.3.6. Why the relaxed criterion suffices for PST

Every load-bearing derivation in the book uses the substrate's spectral-triple structure at finite D: Newton's G via the spectral action $\mathrm{Tr}\, f(D/\Lambda)$ (§8.5); the gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ via the unimodular unitaries of A_F (Chapter 13); gauge bosons and Higgs as inner fluctuations of D (Chapter 13); and the real K-theory (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 (257); full operator equality of bridge$([D_{\mathrm{sub}}, \chi_S])$ with $[D_\alpha, \mathrm{bridge}(\chi_S)]$ is not invoked. Closure under (257) therefore suffices for the Standard Model derivation chain. The inner-fluctuation one-form $A = \sum_i a_i[D, b_i]$ that generates the gauge bosons and Higgs is an operator-level object, but it is built on the *continuum* spectral triple $M \times F$ that equation (257) delivers as its endpoint. It lives over the continuum algebra $C^\infty(M) \otimes A_F$, with A_F fixed at $\mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, and is not part of the substrate-to-continuum bridge data the relaxed criterion controls. The criterion supplies the continuum Dirac operator and its spectral data, on which the (exact, post-limit) inner fluctuations are then constructed in the standard way, so the inner-fluctuation case consumes only the spectral data that (257) preserves.

The distance framework is [116]'s, the same machinery already used to lift L_round closure to quantum Gromov–Hausdorff convergence. Computations 33, 40, 42 establish equation (257) at *finite-N matched truncation*, at Bergman dimension $N \sim 2.5\, D$ (the working truncation of those computations; the asymptotic matched scaling $N(D) = 2\sqrt{D}$ of the refined-bridge lemma chain below is a different convention, chosen for the infinite-N limit, and the two are not to be compared). At the Walsh-weight cutoff $k_D = \lfloor \sqrt{D} \rfloor$, the gap $\gamma_D = |1 - \mathrm{ratio}|$ at $N = 20$ reaches $\gamma_D = 0.007$ at $D = 8$ (Computation 36, offline; the published script runs $N = 10$), within 0.7% of the substrate target, with empirical fit $\gamma_D \sim 19.3 \cdot e^{-0.93D}$ across $D = 4, 6, 8$.

18.3.6.1. 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 (§18.3.5). 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 $k_D = 2 \to 3$ (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-propinquity statement on the infinite Bergman space.

18.3.7. Refined bridge: the Lemma 5–8 closure of L_comm

The infinite-N closure of equation (257) requires a bridge whose substrate-to-Bergman ratio matches $2\sqrt{k}$ uniformly across $k \le k_D$, smoothing the cutoff discontinuity across the Bergman cutoff transitions. Computations 39–42 construct the *symmetric-monomial* bridge $b_C(\chi_S) := T_{f_{|S|}}$ with $f_k(z_0, z_1) = (z_0 z_1)^{m(k)-1} + \beta(k)(z_0 z_1)^{m(k)} + \alpha(k)(z_0 z_1)^{m(k)+1}$, whose gap $\gamma(k)$ is D-independent, eliminating the cutoff discontinuity and making the Mosco averaging trivial. Four lemmas close the construction analytically, and we state each as a result; the full derivations are in §18.4.

Lemma 5 (ratio match; Computations 56–58). Because the bridge symbol depends only on $w = z_0 z_1$, the constrained sup over ∂B^2 reduces to a 1D optimisation, giving the closed-form single-monomial ratio (263), which equals $2 = 2\sqrt{1}$ at $m = 1$ exactly. Since the ratio is strictly increasing in m and asymptotically $\sim m + 1/2$, for every k the target $2\sqrt{k}$ is bracketed by consecutive monomial ratios, and Bolzano's theorem supplies a unique closing $\alpha(k)$ with $R_{m(k)}(\alpha(k)) = 2\sqrt{k}$ exactly (eq. (264)), realised as a rational function of an algebraic root (265).

Lemma 6 (Berezin compatibility; Computation 61). The reproducing-kernel adjoint identity gives $B(T_f T_g^*)(z) = f(z)\overline{g(z)}$ exactly on the infinite Bergman space (eq. (266)), so the refined bridge realises the substrate's weight-class structure as a holomorphic-symbol algebra, up to an $O(|z|^{2N})$ boundary correction that decays exponentially in N.

Lemma 7 (finite-N closure rate; Computations 60, 65). At weight $k = 1$ the gap is closed in elementary form, $R(N) = 2\sqrt{1 - 1/N}$ exactly (eq. (268)), giving leading correction $-1/N$. For $k \geq 2$ the per-k leading constant is ill-conditioned at moderate N (diagnosed in Computations 63, 64), so the foundational bound was established by direct measurement of $|\mathrm{gap}(k, N)| \cdot N \leq 3.7$ across the operative regime, and is now derived at its maximiser and in its large-N limit (Computations 213, 214), $k \leq k_D$ (eq. (269), median 0.95, Computation 150; certified in exact arithmetic over $k \leq 20$ by Computation 203, and the single-monomial part of the problem closed there for every m), yielding the L_comm closure rate $\gamma_D = O(1/\sqrt{D})$ uniformly in $k \leq k_D$ (eq. (270)), polynomial in the substrate size, with a constant that may grow as $D^{1/8}$ beyond the tested range (Chapter 19), which leaves the rate polynomial.

Lemma 8 (spectral-action invariance; Computation 62). At the matched substrate-to-Bergman scaling (271) the eigenvalue counts match (2^D exactly on the substrate side) and the spectral actions are commensurate: $S_{\mathrm{sub}}/2^D = e^{-1}$ exactly, while $S_{\mathrm{Bergman}}/2^D$ converges to a D-independent constant, the standard QGH-convergence statement applied to spectral actions.

The L_comm closure programme is a complete theorem at the relaxed-criterion level, rigorous in all parts except the uniform-in-k bound of Lemma 7(b), which is derived at its maximiser and in its large-N limit and scanned between (Computations 213, 214). A proof over the whole regime is the one remaining quantitative refinement (Chapter 19). The relaxed criterion is sufficient for the downstream spectral-action and gauge-group recoveries (§ *Why the relaxed criterion suffices for PST* above). The closed sub-results are:

- the ratio-matching theorem Lemma 5(a, b, c) (Computations 56–58);
- the Berezin-transform identity Lemma 6 (Computation 61, equation (266));
- the closed-form finite-N gap Lemma 7(a) (Computation 60, equation (268));
- the uniform-in-k direct-gap bound Lemma 7(b) (Computation 65, equation (269)), established by measurement across $k \in [1, 16]$ and $N \in [12, 32]$ there, since extended to $k \leq 80$ and $N \leq 600$ and derived at its maximiser (Computations 203, 214);
- and the spectral-action invariance corollary Lemma 8 (Computation 62, equation (271)).

The foundational closure rests on the direct measurement of $|\mathrm{gap}(k, N)| \cdot N$ across the relevant (k, N) range (Computation 65), giving the L_comm closure rate $\gamma_D = O(1/\sqrt{D})$ uniformly in $k \leq k_D$. See also Appendix B rows 22–65.

18.4. Dirac convergence: the L_round and L_comm lemma chain

This section gives, in full, the two heaviest pieces of the Dirac-operator convergence programme of §18.3: the quantum-Gromov–Hausdorff section round-trip (the L_round side, with its block-decomposition derivation), and the refined-bridge lemma chain (Lemmas 5–8, the L_comm side, with the ratio-matching, Berezin, and finite-N closed forms).

18.4.1. The L_round side: the Gromov–Hausdorff round-trip

Let $\pi_D : A_D \to C(S^3)$ be Walsh-coefficient truncation at weight $|S| \leq k_D := \lfloor \sqrt{D} \rfloor$, and let $\sigma_D : C(S^3) \to A_D$ be the trivial section (the same low-weight content viewed in A_D). The PST analogue of [119] Lemma 3.2 asks for a sequence $\gamma_D \to 0$ and a Lip-norm L_D on A_D with

$$\| T - \sigma_D(\pi_D(T)) \|_{\mathrm{op}} \;\leq\; \gamma_D \cdot L_D(T) \qquad \text{for all } T \in A_D. \tag{258}$$

A numerical sweep across $D = 4, 6, \ldots, 20$ 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 (258) (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 (258) 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

$$L_F(T) \;:=\; \sup_{j\geq 0} \frac{\|\mathrm{ad}_D^j(T)\|_{\text{op}}}{j!}, \qquad \mathrm{ad}_D(T) := [D, T], \;\; D = \sum_a \chi_a, \tag{259}$$

the standard smooth-subalgebra Lip-norm of the spectral triple. Higher-order Dirac commutators of single Walsh modes satisfy $\|\mathrm{ad}_D^k(\chi_S)\|^{1/k} \to 2\sqrt{D}$ 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 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 $\|\mathrm{ad}_D^j(\chi_S)\|_{\text{op}} \leq (2\sqrt{D})^j \|\chi_S\|_{\text{op}}$ uniformly in j, S, D. Combined with Stirling this proves the single-mode L_round closure $\|\chi_S\|_{\text{op}}/L_F(\chi_S) \leq e^{-2\sqrt{D}} \to 0$ as a theorem, not merely a numerical observation. Direct measurement on the constant-coefficient tail (Computation 23) yields $\gamma_D \approx 0.37 \cdot e^{-0.28\,D}$ 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 $j \geq 1$,

$$\|\mathrm{ad}_D^j(T)\|_{\mathrm{op}} \;=\; (2\sqrt{D})^j \cdot \max\big(\|T_{+-}\|_{\mathrm{op}}, \|T_{-+}\|_{\mathrm{op}}\big), \tag{260}$$

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 $\gamma_D = \|T_{\mathrm{tail}}\|_{\mathrm{op}} / (\max(\|T_{+-}\|, \|T_{-+}\|) \cdot M(D))$. Lemma 2 (tail-sum L_F scaling)[1] is thereby *reduced* from the open interference question of higher-order commutators to the explicit operator-norm bound on the structured off-diagonal block $P_+ T_{\mathrm{tail}} P_-$. Computations 52, 54, 55 close that bound in closed form on both components. The operator $P_+ T_{\mathrm{tail}} P_-$ commutes with D_{sub}, hence with the fermionic-signed S_D representation $\rho_F(g)|S\rangle = \mathrm{sgn}(g \mid S)\,|g(S)\rangle$ (the naive bit-permutation rep does *not* commute with D_{sub} because of the Jordan–Wigner σ_z strings). On the ρ_F-trivial component (Computation 52; the underlying algebraic identity $P_+ T_{\mathrm{tail}} = T_{\mathrm{tail}} P_+$ on the projector image is Computation 53) the dominant singular value is $\sigma_1 = (\alpha(D) - \beta(D))/2 = 2^{D-1} - D$ *exactly*, with $\alpha(D) = 2^D - 1 - D - \binom{D}{2}$ and $\beta(D) = -\binom{D-1}{2}$, achieved on the explicit weight-$\{0,1\}$ eigenvector $v_1 = \frac{1}{\sqrt{2}}\big[|0\rangle - \frac{1}{\sqrt{D}} \sum_a |e_a\rangle\big]$ of D_{sub}. On the ρ_F-standard-rep component (Computations 54, 55) the next singular value is $\sigma_2 = D - 2$ *exactly*. By Pieri on $\Lambda^w(\mathbb{C}^D)$ the standard rep $[D-1, 1]$ appears in H_1 and H_2 only. Take the S_{D-1}-fixed vectors $u_1 = |e_0\rangle - (1/D) \sum_a |e_a\rangle$ at weight 1 and $u_2 = -\sum_{b>0} |\{0, b\}\rangle$ at weight 2. On these, the identities $D_{\mathrm{sub}} u_1 = u_2$, $D_{\mathrm{sub}} u_2 = D u_1$ (no weight-3 leakage; the Jordan–Wigner orderings on each triple $\{0, b, c\}$ cancel exactly), $T_{\mathrm{tail}} u_1 = \beta\, u_1$ and $T_{\mathrm{tail}} u_2 = \gamma\, u_2$ with $\gamma = (D-2)(5-D)/2$ collapse T_{tail}^{+-} to the 2×2 rank-1 multiplicity-space matrix $M = \frac{D-2}{2} \binom{1}{1} \begin{pmatrix} -1 & 1 \end{pmatrix}$, whose singular values are $D - 2$ and 0. All other ρ_F-isotypic components vanish identically under T_{tail}^{+-}.

[1] The lemma labels follow the $L_{\mathrm{round}}/L_{\mathrm{comm}}$ closure structure: Lemma 2 and Lemma 2-prime address the L_{round} side, and Lemmas 5(a, b, c), 6, 7(a, b), 8 the L_{comm} side; the non-consecutive numbering reflects steps subsumed into other lemmas or following from [119] Lemma 3.2.

Hence

$$\left\| P_+ \, T_{\text{tail}} \, P_- \right\|_{\text{op}} \;=\; \max(\sigma_1, \sigma_2) \;=\; 2^{D-1} - D \tag{261}$$

exactly for every $D \geq 4$ at $k_D = 2$, with an exponential gap $2^{D-1} - D \gg D - 2$ between the two components. Lemma 2-prime is therefore a theorem, the asymptotic rate is $\gamma_D = O(D^{1/4} e^{-2\sqrt{D}})$ (the pre-asymptotic empirical exponential fit $0.37 \cdot e^{-0.28D}$ of Computation 23 crosses over to the square-root rate near $D \approx 51$), and L_round closes rigorously. Combined with the bridge construction of [116] (Definition 5.1, Theorem 5.2), this gives convergence of $(\mathrm{Lip}(A_D), L_F) \to (C(S^3), L)$ 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 [119] 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 $\|\mathrm{ad}_D^j(T)\|/j!$ 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.

18.4.2. The L_comm side: the refined-bridge lemma chain

18.4.2.1. Refined bridge, symmetric-monomial construction. The infinite-N closure of equation (257) 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 39–42 construct the *symmetric-monomial* bridge $b_C(\chi_S) := T_{f_{|S|}}$ with

$$f_k(z_0, z_1) \;=\; (z_0 z_1)^{m(k)-1} \;+\; \beta(k)\,(z_0 z_1)^{m(k)} \;+\; \alpha(k)\,(z_0 z_1)^{m(k)+1}. \tag{262}$$

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 $k = 1, 2, \ldots, 12$, 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.

18.4.2.2. Refined bridge: the ratio match (Lemma 5). The ratio-match step of the refined-bridge closure, the identification of bridge parameters $(m(k), \alpha(k), \beta(k))$ that give substrate-target ratio $2\sqrt{k}$, is now a closed analytical theorem (Computations 56–58). Since the bridge symbol f_k depends only on $w := z_0 z_1$, $J_z f_k \equiv 0$ and the matrix-valued symbol $M_{f_k} = \sigma_+ J_+ f_k + \sigma_- J_- f_k$ is anti-diagonal in the Pauli basis, so $\|M_{f_k}\|_{\mathrm{op}} = \max(|z_0|^2, |z_1|^2) \cdot |f_k'(w)|$. Writing $r^2 := |z_0|^2$ and taking $r \geq 1/\sqrt{2}$ WLOG by symmetry gives $|w| = r\sqrt{1 - r^2}$, and the constrained sup over ∂B^2 reduces to a 1D optimisation in r.

For the bare single-monomial symbol $f(w) = w^m$, the optimum sits at $r^2 = (m + 1)/(2m)$ and gives the closed-form ratio

$$\mathrm{ratio}(m) \;:=\; \frac{\sup_{\partial B^2} \|M_f\|_{\mathrm{op}}}{\sup_{\partial B^2} |f|} \;=\; m^{1-m}\,(m+1)^{(m+1)/2}\,(m-1)^{(m-1)/2} \quad (263)$$

(Lemma 5(a), Computation 56). At $m = 1$ this evaluates to $2 = 2\sqrt{1}$, the substrate target at weight $k = 1$, so the bare $f_1(w) = w$ closes L_comm exactly at $k = 1$ with zero refinement parameters. Asymptotically $\mathrm{ratio}(m) \sim m + 1/2$, so the single-monomial bridge undershoots the substrate target $2\sqrt{m}$ at $m = 2$ and overshoots for $m \geq 3$; the closing $(m(k), \alpha(k))$ for $k \geq 2$ is supplied by the refinement.

For the 2-monomial family $f_\alpha(w) = w^m + \alpha\, w^{m+1}$, the ratio $R_m(\alpha) := \sup \|M_{f_\alpha}\|_{\mathrm{op}} / \sup |f_\alpha|$ is continuous in α with $R_m(0) = \mathrm{ratio}(m)$ and $\lim_{|\alpha| \to \infty} R_m(\alpha) = \mathrm{ratio}(m + 1)$. Since the closed-form $\mathrm{ratio}(m)$ is strictly

increasing in m, for every $k \geq 1$ the substrate target $2\sqrt{k}$ is bracketed by $\text{ratio}(m(k)) \leq 2\sqrt{k} \leq \text{ratio}(m(k)+1)$ for a unique $m(k) \in \mathbb{N}$ (verified explicitly for $k = 1, \ldots, 12$ and extended to all $k \geq 1$ by the $\text{ratio}(m) \sim m + 1/2$ asymptotics). Bolzano's theorem then gives

$$\exists\, \alpha(k) \in \mathbb{R} \quad \text{such that} \quad R_{m(k)}(\alpha(k)) \;=\; 2\sqrt{k} \quad \text{exactly,} \tag{264}$$

the existence statement of Lemma 5(b) (Computation 57). The closing $\alpha(k)$ admits an explicit parametric expression. Setting $q := 2r^2 - 1 \in (0, 1)$, the critical-point relation for the sup gives

$$\alpha(q) \;=\; \frac{2\,m\,(1 - m\,q)}{(m+1)\,\sqrt{1-q^2}\,((m+1)\,q - 1)}, \qquad q \in \Big(\frac{1}{m+1},\, \frac{1}{m}\Big), \tag{265}$$

and the closing equation $R_{m(k)}(q^*(k)) = 2\sqrt{k}$ is an algebraic equation in $q^*(k)$ of degree $2m(k) + 3$ after rationalising the $\sqrt{1-q^2}$ factor. The unique root $q^*(k)$ in the open interval $(1/(m(k)+1),\, 1/m(k))$ then determines $\alpha(k) = \alpha(q^*(k))$ by (265) (Lemma 5(c), Computation 58). This realises $\alpha(k)$ as a rational function of an algebraic number.

Lemma 5 (parts (a), (b), (c)) thus closes the L_comm *ratio-matching* step of the refined-bridge construction as a complete analytical theorem. The 3-monomial refinements (Computations 41, 42) attain the same ratio match at typically smaller $|\alpha|$ values via the additional β parameter, providing extra flexibility for the downstream structural conditions discussed below.

18.4.2.3. Refined bridge, Berezin compatibility (Lemma 6). The Berezin-transform compatibility of the refined bridge with the weighted-Bergman spectral-triple structure is closed by direct calculation (Computation 61). For holomorphic symbols f, g on B^2 and the canonical weighted Bergman space $H^2_\alpha(B^2)$, the reproducing-kernel adjoint identity $T_g^* k_z = \overline{g(z)}\, k_z$ yields

$$B\big(T_f\, T_g^*\big)(z) \;=\; f(z)\,\overline{g(z)} \qquad \text{exactly on the infinite Bergman space,} \tag{266}$$

for every z in the open ball. Applied to the refined bridge $b_C(\chi_S) = T_{f_{|S|}}$ with holomorphic polynomial symbols $f_k = w^{m(k)} + \alpha(k)\, w^{m(k)+1}$, this is verbatim Lemma 6:

$$B\big(b_C(\chi_S)\, b_C(\chi_T)^*\big)(z) \;=\; f_{|S|}(z)\,\overline{f_{|T|}(z)} \quad \text{exactly.} \tag{267}$$

On the truncated Bergman space $H^2_\alpha(B^2)^{(N)}$ the identity holds up to an $O(|z|^{2N})$ boundary correction that decays exponentially in N for z in the open ball. The bridge therefore realises the substrate's weight-class structure on the Bergman side as a holomorphic-symbol algebra, with the Berezin transform recovering the symbol algebra exactly.

18.4.2.4. Refined bridge: finite-N closure rate (Lemma 7). Computation 59 documents the uniform-in-k closure at finite N. At finite Bergman truncation $H^2_\alpha(B^2)^{(N)}$, the L_comm ratio departs from its infinite-N value $2\sqrt{k}$ by a controlled correction. For the bare single-monomial bridge at weight $k = 1$ ($f_1(w) = w$), the correction is closed in elementary form (Computation 60). In the canonical Bergman weight $\alpha = 0$ with orthonormal monomial basis $e_{a,b} = z_0^a z_1^b/\sqrt{N_{a,b}}$ restricted to $a + b \le N$, the holomorphic Toeplitz operators $T_{z_0^p z_1^q}$ act as diagonal shifts $e_{a,b} \mapsto (\text{coeff}) \cdot e_{a+p,b+q}$ (orthogonal target). Operator norms read off as maximum matrix elements: $\|T_w\|_{\text{op}} = N/[2\sqrt{(N+1)(N+2)}]$ (maximiser $a = b = (N-2)/2$, even N) and $\|T_{z_0^2}\|_{\text{op}} = \sqrt{N(N-1)/((N+1)(N+2))}$ (maximiser $a = N-2$, $b = 0$). The commutator $[D_\alpha, T_w \otimes I_2] = T_{z_0^2} \otimes \sigma_+ + T_{z_1^2} \otimes \sigma_-$ is anti-diagonal in the Pauli basis, giving the closed-form ratio

$$R(N) \;=\; \frac{\|[D_\alpha, T_w \otimes I_2]\|_{\text{op}}}{\|T_w\|_{\text{op}}} \;=\; 2\sqrt{1 - 1/N} \quad \text{exactly (even } N\text{).} \tag{268}$$

Hence $\text{gap}(1, N) := R(N) - 2 = 2\sqrt{1 - 1/N} - 2$ with Taylor expansion $-1/N - 1/(4N^2) - 1/(8N^3) - O(1/N^4)$; the leading coefficient is -1 exactly. At the matched scaling $N(D) \sim 2\sqrt{D}$ this gives $\gamma_D(k = 1) = -1/(2\sqrt{D}) + O(1/D)$ (Lemma 7(a)).

For $k \geq 2$ the 2-monomial bridge $T_{f_k} = T_{w^{m(k)}} + \alpha(k)\, T_{w^{m(k)+1}}$ couples basis vectors through the diagonal shift $(a, b) \to (a+1, b+1)$, and the operator norm requires diagonalising a tridiagonal Toeplitz-like block at finite N. In the bulk-large-a limit the coefficients stabilise to a constant tridiagonal with asymptotic operator norm $(\alpha(k) + 2)/2^{m(k)+1}$; the finite-$N$ correction is again $O(1/N)$ by the classical Bergman–Toeplitz convergence theorem for polynomial symbols ([144], [145]). Per-k extraction of the leading constant $c_1(k)$ via least-squares fit is ill-conditioned at the moderate N reachable by direct dense linear algebra: two-parameter (Computation 63) and three-parameter (Computation 64) fits over $N \leq 34$ disagree by factor 2–10 and frequently in sign for $k \geq 2$, the leading $1/N$ term not yet separated from higher-order structure when N is comparable to the bridge degree $m(k) + 1$. Lemma 7(b) is closed instead by direct measurement of the gap-times-N quantity $G(k, N) := |\mathrm{gap}(k, N)| \cdot N$ itself (Computation 65), bypassing the $c_1(k)$ extraction. Across the operative regime $k \leq k_D$, with $k \leq 20$ and $N \in [12, 40]$ (195 measured (k, N) pairs, Computation 150):

$$|\mathrm{gap}(k, N)| \cdot N \;=\; |R(k, N) - 2\sqrt{k}| \cdot N \;\leq\; 3.7 \quad \text{for } k \leq k_D = N/2, \quad (269)$$

median 0.95, the supremum attained at the *smallest* substrate size in range. **The restriction to $k \leq k_D$ is necessary, not cosmetic.** Computation 150 measures $G = 22.0$ at $(k, N) = (39, 22)$, so the bound does not hold uniformly in (k, N); every violation lies at $k > k_D$, which the Walsh cutoff excludes. Computation 65's reported maximum 5.249 at $(k, N) = (15, 20)$ is itself such a point, since $k_D = 10$ there. Hence at the matched scaling $N(D) = 2\sqrt{D}$,

$$\gamma_D \;=\; \max_{k \leq k_D} |\mathrm{gap}(k, N(D))| \;\leq\; \frac{3.7}{2\sqrt{D}} \;=\; O\!\left(\frac{1}{\sqrt{D}}\right) \quad \text{uniformly in } k \leq k_D \quad (270)$$

(Lemma 7(b)). The foundational L_comm closure rate is therefore $\gamma_D = O(1/\sqrt{D})$, polynomial in the substrate size, weaker than the exponential L_round rate but uniformly controlled. **That rate is tight, not merely an upper bound.** At fixed k the quantity G converges to a nonzero constant

of order one as N grows, at $k = 1$ it runs $1.0161, 1.0128, \ldots, 1.0057$ over $N = 16$ to 44, monotone towards about 1.00, while $m(k)/N$ falls from 0.062 to 0.023, so $\gamma_D = \Theta(1/\sqrt{D})$ and no better (Computation 153). An apparent downward trend in the matched-scaling sequence is sawtooth phase in the $m(k)$ stratum, not decay. This is consistent with the vanishing Englis–Zhu deficit of equation (282): a relative deficit $O(m/N)$ on the norms and an $O(1/N)$ gap in their ratio are different statements, and both hold. The per-k closed form of $c_1(k)$ for $k \geq 2$ is a separate quantitative question: the numerics suggest no simple pattern, but the foundational closure does not require it.

18.4.2.5. Refined bridge: spectral-action invariance (Lemma 8). The Connes spectral-action invariance under the refined bridge b_C follows as a corollary of Lemmas 5, 6, 7(a) (Computation 62). At the matched substrate-to-Bergman scaling

$$\Lambda(D) \;=\; \sqrt{D}, \qquad \ell(D) \;=\; (3 \cdot 2^D)^{1/3}/\Lambda \;=\; 3^{1/3}\, 2^{D/3}/\sqrt{D}, \tag{271}$$

the eigenvalue counts are commensurate: $\#\{\lambda(D_{\text{sub}}) : |\lambda| \leq \Lambda\} = 2^D$ exactly, and $\#\{\lambda(D_\alpha) : |\lambda| \leq \Lambda\} \approx \frac{4}{3}(\ell\Lambda)^3 = 4 \cdot 2^D$ with the multiplicity $2(n+1)(n+2)$ counted for both signs (the level count without multiplicity is $(\ell\Lambda)^3/3 = 2^D$; the Bergman count up to the cutoff scales as the substrate Hilbert-space dimension, with a D-independent factor of four). At this scaling, the spectral actions are commensurate:

$$\begin{aligned} S_{\text{sub}}(\Lambda, f) \;/\; 2^D \;&\to\; \text{const}_{\text{sub}}(f), \\ S_{\text{Bergman}}(\Lambda, \ell, f) \;/\; 2^D \;&\to\; \text{const}_{\text{Bergman}}(\ell, f) \quad \text{as } D \to \infty, \end{aligned} \tag{272}$$

both as bounded $O(1)$ quantities independent of D. Numerical verification with Gaussian cutoff $f(x) = e^{-x^2}$ across $D = 6, 8, 10, 12$: $S_{\text{sub}}/2^D = e^{-1}$ exactly; $S_{\text{Bergman}}/2^D$ converges to a D-independent constant (≈ 21.7 at the $j = 1$ moment, ≈ 54.2 at $j = 2$, ≈ 189.7 at $j = 3$). Exact equality requires the standard Chamseddine–Connes choice of cutoff function f (absorbing the ratio into a renormalisation of f). This is the standard QGH-convergence

statement applied to spectral actions.

18.5. Summary

This chapter carries the proofs the physics chapters import. The Mosco-convergence theorem is stated with its seven assumptions A1–A7, the structural ones verified and the equidistribution condition A3 supplied by P3's maximum-entropy sublimation (§18.1). The K-uniform rate is proved for the i.i.d. surrogate by a five-step Berry–Esseen argument with three closures, from which equicoercivity A7 follows by the reduction lemma (§18.1.6). Dirac-operator convergence, which the spectral-triple framework requires beyond the Laplacian's, is reduced to a single analytic estimate (§18.3) and carried through the L_round round-trip and the L_comm lemma chain under the relaxed norm-equivalence criterion (§18.4); the uniform bound's sharp value and its limit law are derived in Chapter 19's account of the frontier. No reader needs this chapter to follow the physics; the rest of the book needs it to exist.

19. Resolved Problems and the Open Frontier

A theory should say plainly what it has settled and what it has not. This chapter does both: first the problems PST resolves, each pointing back to where it was derived rather than repeating the derivation, and then the frontier that is still open. The second list is the more important one, and it is stated in full rather than softened.

This chapter takes stock of the research problems PST resolves and states, in one place, the frontier that remains open. The resolved entries are deliberately brief: each is an index entry for a result derived at length earlier in the book, and says where to find it rather than repeating it. The open and reduced items that follow are correspondingly long, because an open scope has to be stated exactly. The resolved results come first. They open with the four sectors (the foundational object and the emergent geometry, gravity and cosmology, quantum mechanics, and the Standard Model) and then the harder individual closures that underlie them. PST's standing among substrate theories follows, and the open frontier is collected at the end (§19.9). Each resolved result below is established in the body and supported by an explicit verification; each open item is stated as an honest bound, not a closure.

19.1. The resolved sectors

The four sectors come first, each as a single entry: the foundational object and the emergent geometry, gravity and cosmology, quantum mechanics, and the Standard Model. Each says what is settled and points at where it was derived, rather than deriving it again; the harder individual closures that underlie them follow at length.

19.1.1. The foundational object

The foundational locus of PST (§1.5) is a real spectral triple of total real K-theory (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 (mod 8), not a count of spatial dimensions and not the critical dimension of superstring theory. The construction is verified explicitly (Computation 3): 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.

19.1.2. Emergent geometry and the Lorentzian signature

The passage from substrate to spacetime is carried by the modal-sublimation projection functor Φ, and it is essentially unique: Φ is fixed up to diffeomorphism, the residual freedom in the realisation map ρ being characterised exactly by diffeomorphism invariance and nothing more (§4.11). Diffeomorphism invariance of the emergent geometry is therefore not imposed but derived; it is the precise statement of the remaining ambiguity in Φ. The metric *signature* is likewise derived rather than assumed: the Lorentzian $(+,-,-,-)$ is forced as a theorem by hyperbolic well-posedness of the Goldstone wave equation along the global bias $\hat{\tau}$, which admits a conserved Noether current only for signature $(1, n-1)$ (§4.12, Computation 2). PST thus recovers the Lorentzian signature, rather than postulating it, as a structural consequence of the tension geometry.

19.1.3. Gravity

The Einstein–Hilbert action is the leading term of the spectral action $\mathrm{Tr}\, f(D/\Lambda)$, with Newton's constant entering as a consistency relation among (m_h, M_P, v) rather than a free input (Computation 5). 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 4). The cosmological term completes the gravitational sector: its sign and equation of state $w = -1$ are structural theorems (§16.5), its magnitude $T(C)$-contingent (Computation 82). The substrate coherence length d_0 further predicts a

Casimir correction of derived form, falsifiable through its exponent, though at $d_0 = 2.2$ am its amplitude is far below present reach. The strong-field limit is treated in §8.3: the projected field equations cover the horizon with some forty-two orders of margin, so the thermodynamic entropy $S = A/4G$ follows from them by Wald's Noether-charge theorem (§8.3.3, Computation 154). The register drops in two remaining places: the horizon's substrate-side reading is an *identification* rather than a theorem, and the statistical entropy and what replaces the central singularity are open (§19.9.3).

19.1.4. Cosmology

The Cosmological Principle, spatial homogeneity and isotropy, is a theorem rather than a postulate: it follows from the permutation invariance of the Bernoulli measure μ under finite relabelling of the distinction index, which renders the instantiation rate j_τ spatially uniform and isotropic at all scales large relative to d_0 (§16.2). The Friedmann–Robertson–Walker metric is then the unique homogeneous isotropic solution of the emergent Einstein equations, so the Friedmann equations are derived, not assumed, and cosmic acceleration arises from ongoing sublimation rather than a tuned constant (Chapter 16). The horizon and flatness problems dissolve without inflation: shared early-universe temperature and spatial flatness follow from origin-free simultaneous instantiation under permutation invariance, not from any signal propagating through a pre-existing space (Chapter 16). The cosmological-constant sign and $w = -1$ enter here as the structural theorems already noted under Gravity (§16.5, Computation 82).

19.1.5. Quantum mechanics

The canonical commutation relations follow from the gradient structure on configuration space, via Mosco convergence of the Boolean Dirichlet forms (§10.6: $\hat{p}_\mu = -i\hbar\,\Pi(\partial^{\mathrm{B}}_\mu)$). The differentiable structure is supplied by the Boolean gradient derived from the symmetric-difference metric, and the Bernoulli measure $\mu = \bigotimes_{a\in D} \mathrm{Bern}(1/2)$ is uniquely fixed by relabelling and per-property complementation invariance, with no additional primitive.

The Born rule is a theorem here, not a postulate: the identification of $|\psi|^2$ with measurement probability follows from the Bernoulli measure and Gleason's theorem applied to $L^2(\mathcal{P}(D), \mu)$, given the von Neumann coupling and the unitary propagation between crossings that PST inherits (§10.2). The projective content, the collapse of the wave function onto an eigenspace, is derived from threshold-crossing of the modal tension (§10.3). Both contents of the measurement axiom are thus recovered at proof-detail level; what remains is textbook axiomatic-field-theory work and the one regime question of §10.3.2.5, the coupling at which the einselected pointer basis rather than (98) does the separating.

19.1.6. The Standard Model

The gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ 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 6). The colour $\mathrm{SU}(3)_c$ is the $M_3(\mathbb{C})$ factor; one generation's colour and charge assignments follow from the octonionic origin of A_F. The chiral $\mathrm{SU}(2)_L$ acts chirally through the Connes order-one condition: the chiral bimodule support of $\mathbb{H} \subset A_F$ on $H_F^{(+)}$ (and not $H_F^{(-)}$) is structurally derived from precausal postulates via Computations 7, 8, and 11 together. Computation 8 shows that γ_M does not furnish the chirality but the substrate's $\mathrm{so}(4) = \mathrm{su}(2)_L \oplus \mathrm{su}(2)_R$ isometry of S^3 does; Computation 11 verifies both parity-conjugate representations have clean block support; the directed modal threshold of P3 selects the parity-violating configuration. Given the bimodule, the order-one condition selects the Yukawa D_F uniquely (Computation 7).

The Higgs sector matching: the PST effective field theory below the matched scale $M_* = 4\pi m_h \sqrt{2/3}$ equals the Standard Model exactly (§12.2, equation (127)). The substrate-side factor of the field-normalisation ratio $Z^2 = \lambda_{\mathrm{SM}}(M_*)/b(M_*) \approx \mathrm{e}^{-1}$ is derived structurally from P1–P3, and bridge premise B is reduced within PST to a field-strength-renormalisation identity. What it rests on is premise F4, that the Standard Model (SM) is PST's effective field theory (EFT) below M_* (§12.6.2.1), together with the renormalisation-group

equation (RGE) test (§19.7).

Two further Standard-Model consistency results complete the sector. Under $N_c = 3$, the dimension of the $M_3(\mathbb{C})$ colour factor of A_F, the SM hypercharge pattern of eq. (194) is anomaly-free and follows from the standard determination (Gell-Mann–Nishijima, photon masslessness, Yukawa gauge-invariance, and anomaly cancellation) with no free hypercharge parameter beyond the empirical normalisation; equivalently, given the Standard-Model hypercharges, anomaly freedom forces $N_c = 3$ (Chapter 13, Computation 136). As stated in the body, this is a structural consistency theorem tying colour rank to anomaly freedom, a demonstration of mutual consistency rather than a derivation of either the hypercharges or anomaly cancellation from the postulates alone.

19.2. 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, L_round and L_comm, under the relaxed norm-equivalence criterion that suffices for PST. The full proof, with all lemmas and the relaxed-versus-strict discussion, is §18.3; the supporting computations are summarised here.

The section round-trip (L_round) is a quantum Gromov–Hausdorff (QGH) statement, giving QGH convergence of the Walsh substrate algebras to $C(S^3)$ (Computations 20–23). The single-mode closure is a theorem (Computation 50): the Jordan–Wigner Clifford structure gives $D^2_{\text{sub}} = D \cdot I$, so $\|\chi_S\|_{\text{op}}/L_F(\chi_S) \leq e^{-2\sqrt{D}} \to 0$ rigorously by Stirling. The constant-coefficient tail closes in closed form, via the block-decomposition identity (Computation 51) and the fermionic-signed S_D representation (Computations 52, 54, 55). This gives $\|P_+ T_{\text{tail}} P_-\|_{\text{op}} = 2^{D-1} - D$ exactly for every $D \geq 4$ at $k_D = 2$ (equation (261)), so the tail-sum closure is also a theorem; the true asymptotic rate is $\gamma_D = O(D^{1/4} e^{-2\sqrt{D}})$.

The Dirac-aware lift (L_comm) closes under the relaxed criterion (equation (257)) via the refined symmetric-monomial bridge (Computations 39–42).

It is a rigorous theorem across five sub-results, the ratio match, Berezin compatibility, the closed-form $k = 1$ gap, the direct-gap bound $|\mathrm{gap}(k, N)| \cdot N \leq 3.7$ on $k \leq k_D$ and spectral-action invariance at matched scaling, each stated with its computations in §18.3 (Lemmas 5 to 8). 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.

Relaxed criterion: strict-criterion failure and the constant. The strict-criterion homomorphism gap fails at finite (D, N) (Computation 35), which is why the relaxed norm-equivalence criterion is the operative one here. Three refinements sharpen the L_comm closure: a closed form for the Toeplitz symbol, an analytical account of why the bound exists on $k \leq k_D$, and a reduction that removes the Dirac operator from the problem and solves one case of it outright (§19.2.1).

19.2.1. Reducing the bound to ladder commutators

The constant has resisted partly because of what it is a constant *of*: a singular value of a matrix twice the size of the truncated Bergman space, built from the Dirac operator. It need not be. Computation 157 takes the Dirac operator out of the problem.

T_{w^m} shifts $(a, b) \mapsto (a + m,\, b + m)$ and therefore preserves $a - b$, which is twice the J_z eigenvalue: $[J_z, T] = 0$ exactly, for every bridge. What survives in $D_\alpha = J_1 \otimes \sigma_x + J_2 \otimes \sigma_y + J_z \otimes \sigma_z$ is $[J_+, T] \otimes \sigma_- + [J_-, T] \otimes \sigma_+$, and since $\sigma_-^\dagger \sigma_-$ and $\sigma_+^\dagger \sigma_+$ are the two orthogonal projectors while the cross terms vanish, its square is block diagonal. Hence

$$\big\|[D_\alpha,\ T \otimes I_2]\big\| \ = \ \max\big(\|[J_+, T]\|,\ \|[J_-, T]\|\big), \tag{273}$$

and the $z_0 \leftrightarrow z_1$ symmetry of the Bergman weight makes the two equal. The ladder commutator is itself a shift with an integer coefficient: $[J_+, T_{w^m}] = m\, S_{m+1,\, m-1}$ exactly, where $S_{p,q}$ is the normalised shift $(a, b) \mapsto (a + p,\, b + q)$.

Both identities hold to machine precision.

Everything in sight is therefore built from shifts. The bridge $T_{f_k} = T_{w^m} + \alpha(k)\, T_{w^{m+1}}$ and its commutator are each a sum of two shifts whose displacements differ by $(1,1)$, so each is bidiagonal along the chains $(a,b) \mapsto (a+1,\, b+1)$, and its norm is a maximum over many small bidiagonal problems rather than one large one. The resulting evaluator reproduces the brute-force singular value to 10^{-11} and costs $O(N^2)$.

One case then closes in closed form. At $k = 1$ the bridge is the bare monomial ($m = 1$, $\alpha = 0$), so both operators are single shifts, and a shift's norm is simply the largest of its weights. With $\|z_0^a z_1^b\|^2 = 2\, a!\, b!/(a+b+2)!$ the two maxima are attained at opposite corners of the admissible lattice ($\|T\|^2$ at $a = b$, and $\|[J_+, T]\|^2$ at $b = 0$) giving

$$\begin{aligned} \operatorname{ratio}(1,N)^2 \;&=\; \frac{4(N-1)}{N}, \qquad \text{hence} \\ C(1) \;:=\; \lim_{N\to\infty} \big(2 - \operatorname{ratio}(1,N)\big)\, N \;&=\; 1 \end{aligned} \tag{274}$$

exactly. That is derived from the weights rather than fitted, and it matches the measured value to 10^{-12} up to $N = 50$.

The sharp constant survives a fourfold extension of the check. Computation 150 swept $k \le 40$ at $N \le 40$, so the operative regime $k \le N/2$ was in fact tested only to $k = 20$. The reduced evaluator reaches $k = 80$, and the maximum of $G(k,N)$ over the operative regime is still 3.6736, at $(k,N) = (6,12)$, the smallest N admissible for $k = 6$, with the per-k maxima falling to between 1.5 and 2.2 by $k = 50$ to 80. The constant is a boundary value at small N, and it is not drifting upward.

The sharp value is an explicit algebraic number, and its corner is the maximum for a reason. At $(k,N) = (6,12)$ the headroom $N - 2m$ is four, so every chain has at most three points, and both extremal chains are the diagonal $(0,0), (1,1), (2,2)$ (Computation 214). Their 3×3 matrices have exact rational

weights and the characteristic polynomials

$$\begin{aligned} 122367445200\,L^3 - 24820950\,\alpha^2 L^2 - 258112008\,L^2 \\ + 1050\,\alpha^4 L + 13419\,\alpha^2 L \\ + 164108\,L - 30 = 0, \\ 2913510600\,L^3 - 17382750\,\alpha^2 L^2 - 117070800\,L^2 \\ + 21875\,\alpha^4 L + 182000\,\alpha^2 L \\ + 1432256\,L - 5120 = 0, \end{aligned} \tag{275}$$

for $\|T\|^2$ and $\|[J_+, T]\|^2$, with $\alpha = \alpha(6)$ the root in $(1, 2)$ of a degree-24 integer polynomial that Lemma 5(c)'s closure reduces to. So $G(6, 12) = 12\,(2\sqrt{6} - \sqrt{L_C/L_T}) = 3.673556413812905\ldots$, an algebraic number rather than a float, inside Computation 203's enclosure. Why that corner: at fixed headroom h and degree m the value is single-humped in α, with the hump where the two monomials are comparable and its top rising with m, while the operative corner $N = 2k$ forces $h = 2k - 2m(k)$ to grow with k. The integer lattice $(k, m(k), \alpha(k))$ samples those curves at few points, and $k = 6$ is the only one with $h \le 4$ and α on a hump; $k = 5$ shares the headroom with $\alpha = 28.7$, far down the tail. Rescanned over every chain for $k \le 80$ and $N \le 600$ nothing else exceeds 3.1, and Computation 213's limits bound the rest. No analytic mechanism beyond the lattice sets 3.6736; it is the one point of the operative regime where a small headroom meets a balanced bridge.

It is also not the same number as the rate's constant, and the two should not be conflated. Near $N = 250$ the sawtooth narrows to a band and $C(k)$ sits well under one: 1.0011 at $k = 1$, which is $1 + 1/(4N)$, the next order of (274), then 0.83, 0.10, 0.78 and 0.79 at $k = 2, 3, 4, 6$. $C(k)$ is what the rate $\gamma_D = \Theta(1/\sqrt{D})$ carries (Computation 153); 3.674 is the worst finite-N boundary value. Reading the second as the first would overstate the asymptotic constant by roughly a factor of four.

The discrete maximum over a, for every m. Computation 203 supplies the

closed form that the corner case left wanting. Along a fixed degree $s = a + b$ the commutator weight obeys the exact ratio test

$$\frac{f(a+1)}{f(a)} \;=\; \frac{(a+m+2)\,b}{(a+1)(b+m-1)}, \tag{276}$$

which is at least one exactly while $(m+1)\,b \geq (m-1)(a+1)$, and both shift weights increase with the degree, so the maxima sit on the top admissible degree $s = N - 2m$: $\|T_{w^m}\|^2$ at $a = b = s/2$ and $\|[J_+, T_{w^m}]\|^2$ at $a^* = \lfloor((m+1)s - (m-1))/(2m)\rfloor + 1$, $b^* = s - a^*$, for $m \geq 2$ (at $m = 1$ the maximiser is the corner $b^* = 0$, below). The ceiling in a^* *is* the sawtooth. Verified in exact rational arithmetic against every lattice point for $m \leq 8$. The single-monomial finite-N ratio is therefore an explicit factorial ratio for every m, generalising (274); its limit is Lemma 5's closed form, and the deficit $(\mathrm{ratio}_{\mathrm{cf}}(m) - \mathrm{ratio}_m(N))\,N$ tends to 1 at $m = 1$ and to 0 for every $m \geq 2$. That is a theorem, not an observation (Computation 204): both shifts move the degree by $2m$, so the degree-dependent part of the weights cancels in the ratio and what is left is a finite product, whose expansion in $1/x$, $x = s/2$, is a finite Taylor expansion. Writing $a^* = (m+1)x/m + \delta$ with δ the offset of the lattice maximiser from the continuum point, a function of $s \bmod 2m$ alone, the first-order terms cancel identically and

$$\log \mathrm{ratio}_m(N)^2 \;=\; \log \mathrm{ratio}_{\mathrm{cf}}(m)^2 \;-\; \frac{m\,(6m^2\delta^2 - 6m\delta + 1)}{6\,(m^2-1)\,x^2} \;+\; O(x^{-3}), \tag{277}$$

for every $m \geq 2$, so the single-monomial ratio converges at rate N^{-2} with an oscillating coefficient in closed form, $6(m\delta - \frac{1}{2})^2 - \frac{1}{2}$ changing sign along the sawtooth. The cancellation is first-order stationarity at an interior maximum: at $m = 1$ the maximiser is the corner $b^* = 0$, there is no stationarity, and the $1/N$ term survives with coefficient 1.

The two-monomial constant is a limit, and it is in closed form. The bridge $T_{w^m} + \alpha(k)\,T_{w^{m+1}}$ couples the chains, and its $1/N$ term does not cancel: the deficit times N is 0.85, 0.17, 0.82 and 0.81 for $k = 2, 3, 4, 6$ at $N = 480$,

with the extremal chains identified (the diagonal $a = b$ for $\|T\|$; for the commutator a chain whose start drifts with N). Followed to $N = 15360$ by Computation 206 it rises slowly and monotonically, by 0.16 at $k = 2$ over five doublings with no plateau in sight, and for a time the data could not decide between a slowly approached limit and a logarithm. Computation 213 decides it. The deficit times N is the sum of three terms, the weights' $1/N$ term and the two extremal chains' zero-point energies below their symbols, and each has its own law. The weights' term converges to a closed form from the $1/s$ expansion of the Bergman weight. The commutator's chain runs through the corner (r^*, N) of the lattice, r^* the interior maximum of its symbol in the slope $a/(a+b)$ (the position the drifting chain was heading for), and its eigenvector is a harmonic ground state whose zero point is a constant in closed form. The bridge's diagonal chain, whose weights rise monotonically to the truncation, carries an Airy function pressed against the hard edge, whose zero point decays as $A(k)\,N^{-1/3}$. So

$$\big(2\sqrt{k} - \mathrm{ratio}(k,N)\big)\,N \;=\; D_\infty(k) \;-\; A(k)\,N^{-1/3} \;+\; E_N, \qquad (278)$$

with the wall term $E_N = O(N^{-1/2}\sqrt{\log N})$, $D_\infty(2) = 1.2263$ and $D_\infty(6) = 1.5235$, approached from below, which is why no plateau was visible: at $N = 61440$ the $k = 2$ value 1.072 is still 0.15 short of its limit. The Airy edge the problem was expected to need is real, but it sits on the bridge's own norm, not on the commutator's. Over k the limit is below 2.47 for $k \le 80$ and below the certified 3.6736 for every $k \le 482$; its envelope grows slowly, like $k^{0.22}$, and crosses 3.6736 at $k = 483$, at substrate sizes no computation reaches. The sharp value itself is settled below: each chain norm is the top eigenvalue of a nonnegative tridiagonal matrix, so Collatz–Wielandt bounds with a rational test vector enclose it in exact arithmetic, and over $k \le 20$ at the ten smallest admissible N every $G(k,N)$ is certified below 3.6736, the enclosure at $(6,12)$ being narrower than 10^{-7} (Computation 203); at $(6,12)$ the value is an explicit algebraic number (Computation 214, equation (275)). The 3.7 of (269) is derived at its maximiser and scanned on the rest of the

range; every value to $N = 61440$ lies below it, and it holds as $N \to \infty$ for every fixed $k \le 482$, so the closure rate of Lemma 7 is established without the logarithm that was feared. What the limit's slow growth in k costs is a constant in Lemma 7 that may grow like $k_D^{1/4} = D^{1/8}$ far beyond the range tested, which leaves the rate polynomial.

19.2.2. Asymptotic Toeplitz-symbol closed form for $c_1(k)$

The asymptotic operator norm of the refined-bridge symbol $f_k(w) = w^{m(k)} + \alpha(k)\, w^{m(k)+1}$ on the infinite Bergman space is given by the L^∞ norm of the Toeplitz symbol

$$g_k(\theta) \;=\; \frac{1}{2^{2m(k)}}\, \mathrm{e}^{2im(k)\theta}\, \big(1 + \tfrac{\alpha(k)}{2}\, \mathrm{e}^{i\theta}\big), \tag{279}$$

maximised at $\theta = 0$ giving $\|T_{f_k}\|_{\mathrm{op},\infty} = \big(1 + \alpha(k)/2\big)/2^{2m(k)}$. The commutator $\|[D_\alpha, T_{f_k} \otimes I_2]\|_{\mathrm{op},\infty} = 2\sqrt{k} \cdot \|T_{f_k}\|_{\mathrm{op},\infty}$ in the bulk by Lemma 5(c)'s ratio match, and the finite-N correction at leading order $1/N$ is $c_1(k) = O(1)$ for $k \le k_D$ by the bound $|c_1(k)| \le 3.7$ of Computation 150, and derived at its maximiser (Computation 214). Together these constitute a closed form for $c_1(k)$ at proof-detail level: the asymptotic norm is the symbol L^∞ value exactly, and the $1/N$ correction is bounded uniformly; its k-profile is $D_\infty(k)$, in closed form (Computation 213). The foundational closure of Lemma 7(b), and through it the L_comm closure programme as a whole, is delivered by the Toeplitz-symbol asymptotic + uniform-bound combination.

19.2.3. Analytical uniform-in-(k, N) bound

The bound $G(k, N) := |\mathrm{gap}(k, N)| \cdot N \le 3.7$ on the operative regime $k \le k_D$ (Computation 150; derived at its maximiser by Computation 214) is explained in its order of magnitude by the standard Bergman-truncation theory for polynomial-degree Toeplitz operators. It does *not* extend to all (k, N): Computation 150 measures $G = 22.0$ at $(k, N) = (39, 22)$. Specifically, for the truncated Bergman space $H^2_\alpha(B^2)^{(N)}$, the operator-norm deficit of T_{f_k} relative to its infinite-Bergman value is controlled by the projection onto the

boundary layer of total degree in $[N - m(k), N]$, giving

$$\|T_{f_k}\|_{\mathrm{op},N} \;=\; \|T_{f_k}\|_{\mathrm{op},\infty}\big(1 + O\big(m(k)/N\big)\big) \tag{280}$$

uniformly in k by the classical Bergman-truncation convergence theorem for polynomial symbols (Engliš [144], Zhu [145]). The same uniform $O(m(k)/N)$ bound holds for the commutator $[D_\alpha, T_{f_k} \otimes I_2]$ since D_α acts as a bounded operator on $H^2_\alpha(B^2) \otimes \mathbb{C}^2$ and preserves the polynomial-degree filtration up to a shift by 1. The deficit is therefore governed by the bridge degree relative to the truncation, and not by k or N separately. Computation 150 confirms this: bucketed by $m(k)/N$, the measured G stays below 2.7 for $m(k)/N \le 0.25$ and rises sharply as $m(k)/N \to 1/2$. Computation 153 sharpens the variable to the *truncation headroom*

$$s \;=\; N - 2\big(m(k) + 1\big), \tag{281}$$

the room the top bridge monomial w^{m+1} has inside the degree-N truncation; $m(k)/N$ is a proxy that works because $m(k)/N \to 1/2$ is the same event as $s \to 0$. Off the operative regime G is not merely large but *unbounded* as s closes (at $N = 40$ and $s = -2$ it runs $3.9, 7.8, \ldots, 38.1$ across ten consecutive k) so no maximum should be quoted there. What is true is that G is bounded on $k \le k_D$ and unbounded off it, with s marking the boundary.

Why the matched scaling is safe. An earlier version of this argument bounded $m(k) \le \sqrt{D}$ and concluded a deficit $\approx 1/2$, hence $O(1)$ rather than $o(1)$. That bound is the one on k, not on m: the bridge degree satisfies $m(k) \le 2\sqrt{k}$ (Computation 150 verifies this for every $k \le 2000$, with $m(k)/2\sqrt{k} \to 1$ from below). With $k \le k_D = \lfloor \sqrt{D} \rfloor$ and $N(D) = 2\sqrt{D}$ this gives

$$\frac{m(k_D)}{N(D)} \;\le\; \frac{2\sqrt{k_D}}{2\sqrt{D}} \;\le\; D^{-1/4}, \tag{282}$$

verified at every point of the matched-scaling sequence (Computation 150). The deficit therefore *vanishes* at rate $D^{-1/4}$ rather than staying $O(1)$: the

regime $m(k)/N \to 1/2$ in which the estimate degrades is approached only at the smallest substrate sizes and is left behind as D grows. This is the structural reason a uniform bound exists on $k \le k_D$, and it is the content the earlier $O(1)$ reading obscured.

What remains an analytical theorem, rather than a measurement, is the *existence* of a bound $G(k, N) \le C$ on $k \le k_D$ with *some* absolute constant C: the Engliš–Zhu polynomial-symbol convergence theorem supplies a finite C together with the $O(m(k)/N)$ form, and equation (282) supplies the rate at which that argument is applied, but neither delivers the value. The sharp value $C \le 3.7$ was measured (Computation 150, 195 points on the operative regime) and is now derived at its maximiser (Computation 214); the k-profile of the limit is in closed form (Computation 213). **The superseded figure was** $5.249 \le 5.3$ (Computation 65), measured at $(k, N) = (15, 20)$ where $k_D = 10$, outside the regime the cutoff admits, which is why it was both too large and not a bound on anything the theory uses.

19.3. Octonionic-SU(3) structure inside Cl(0, 6)

The octonionic gauge structure is closed at the rigorous level on the $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ sector. §13.5 gives the Casimir-discriminated SU(3) decomposition of $M_8(\mathbb{C}) = 6 \cdot \mathbf{1} + 5 \cdot \mathbf{3} + 5 \cdot \overline{\mathbf{3}} + 2 \cdot \mathbf{8} + \mathbf{6} + \overline{\mathbf{6}}$ to machine precision (Computation 17); the embedding choice is justified by Spin(6)-invariance (§13.7, 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 $\mathrm{Cl}(0, 6) \otimes \mathbb{C}$, so the chirality projector $f = (I + i\omega)/2$ is the unique non-trivial Spin(6)-invariant idempotent, fixing the embedding up to chiral swap; its canonical rank-1 refinement is Furey's primitive idempotent. Furey 2014 [121]'s analysis of the chain algebra generated by left multiplication on $\mathbb{C} \otimes \mathbb{O}$, which is $\mathrm{Cl}(0, 6)$, delivers $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation of fermion content (Computations 3, 19). The $\mathbb{H}$ factor of A_F is supplied separately by the emergent 4-d Lorentzian Dirac sector $\mathrm{Cl}(1, 3) \cong M_2(\mathbb{H})$ via right-$\mathbb{H}$- multiplication on the spinor module (Computation 66), and the synthesis to full $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ goes

through Dixon's hyperspinor framework $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ [146]. The three-generation count, derived from the directional content of V_7 (§14.2), places the SM generations in Furey's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0,6)$ (this is Furey's identification, the standard reading in the octonionic SM programme; the colour-versus-generation distinguishability question [121] flags as open is settled here in the negative, Computation 175: the three sectors are orthogonal, exactly, but orthogonality is not invariance, and $\hat{\tau}^{\perp}$ is the irreducible **3** by Schur, so PST derives one three, not two). Combinatorially, the three quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$ are the natural three-fold partition (§14.2, Computations 46–48).

19.4. Matter, and the theory's empirical exposure

The entries above are index entries, deliberately brief. This one is not, because it does a job nothing else in the book does: it collects, in one place, every point at which PST touches a measurement. A reader deciding whether the theory is worth taking seriously needs that list to be short, specific and honest about which items are predictions, which are consistency checks that any theory reproducing general relativity would also pass, and which are not predictions at all.

19.4.1. Boolean occupation makes spin-statistics structural

The Boolean $\{0,1\}$ distinctions of P1 *are* fermionic occupation numbers. Nothing is added to make them so: the configuration space $\mathcal{P}(D)$ carries a canonical anticommutation-relations algebra because a site is either in a configuration or not, and exclusion is the set-theoretic fact that it cannot be in it twice. The statistics is therefore structural rather than imposed, the spin label being the spinor bundle's (§11.3), (§11.2, §11.3), and it is closed in three directions, two of them given A1 to A4 (Mosco convergence of the canonical anticommutation relations [CAR], the Wightman axioms, and the Dirac equation) rather than asserted in one.

This is worth stating plainly because it is the point at which PST's primitive pays for itself. A theory that began with fields would have to postulate the

spin-statistics connection or import it from relativistic invariance; a theory that begins with distinction gets it from the meaning of the primitive.

19.4.2. A 6-sphere vacuum quantises angular momentum

The instantiated vacuum is a 6-sphere of degenerate directions in V_7 (§4.1). Its 2-plane projection is the familiar ring of minima, so orbital motion is the only stable ground state available: a state at rest at the centre is not a minimum. The angular momentum such motion carries is the Noether charge of the broken U(1), and lies in $\hbar\mathbb{Z}$ by single-valuedness around the ring, the integer coming from the vacuum topology and the unit from canonical quantisation (§9.4, §9.5). Quantisation here is a statement about the shape of the vacuum, not a rule applied to it.

19.4.3. The three places the theory can be measured

One prediction, one scale, and one agreement. They are of different kinds and the difference matters more than the count.

(1) The Casimir correction: exact in form, its size fixed by the normalisation premise. PST predicts a fractional deviation from the standard Casimir force of $\delta F/F = \xi\,(d_0/d)^2$, a d^{-6} component of the force, with the coefficient $\xi = 90/\pi^2 \approx 9.12$ *parameter-free* under the Gaussian projection-kernel choice (Chapter 15, Computations 68, 73). The amplitude is ξd_0^2, and d_0 is fixed by the gradient normalisation: under the canonical normalisation on which M_* rests, $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am (Computation 201), nine orders below the diagnostic bound $d_0 \lesssim 2.3$ nm the data allow, and the correction is 10^{-19} of the force at 20 nm. Table 1 gives the grid at the bound, which is the sensitivity a measurement would have, not the signal the theory expects.

What makes this a test rather than a fitting exercise is the exponent. The dominant systematics carry distinct functional forms (finite-conductivity Lifshitz corrections are d^{-1}-like, roughness goes as σ^2/d^2, thermal as $k_B T d/\hbar c$) so the two-separation ratio test (equation (218)) isolates the d^{-6} component apparatus-independently. A measured d^{-6} component at nanometre amplitude

would refute the canonical normalisation, and with it M_* and the Higgs match; its absence at every reachable sensitivity is what the theory expects. **What the theory cannot survive** is a deviation of the wrong exponent, or a d^{-6} component with the wrong coefficient at an amplitude the normalisation allows.

(2) The modal scale M_*: a derived number that survives the collider data in one reading of what it does, and predicts there. $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV is a parameter-free one-loop derivation (§15.2, Computations 93, 97), and its associated length is the sub-fermi $\ell_* = \hbar c/M_* \approx 1.535 \times 10^{-19}$ m.

Two qualifications travel with it, and omitting either would overstate the result. First, it is *irreducibly a scalar-sector* prediction: the full-SM extension of the boson-fermion cancellation that produces it is settled in the *negative* (Computation 120), which is a definitive limitation on the result's scope rather than a gap awaiting work. Second, the obvious collider bounds do not apply. PST's additional structure is a finite noncommutative internal factor, not large extra spatial dimensions. The Arkani-Hamed–Dimopoulos–Dvali (ADD) limits, Large Hadron Collider (LHC) mono-jet and mono-photon at $M_* \gtrsim 5$–8 TeV, SN1987A and neutron-star Kaluza–Klein (KK) graviton emission, are therefore derived under premises PST does not share, and do not transfer in their standard form (§15.10, Computation 10). The confrontation with PST's own coupling structure is Computation 205: read as a cutoff on the momenta of the emergent quanta, or as a kernel on every vertex, M_* is excluded by the jet and dilepton spectra by enormous margins; read as the projection's kernel on the Higgs doublet alone, which is what §15.2 states, it is consistent with every Higgs measurement and predicts a deficit in the Higgs transverse-momentum spectrum of 23 percent above 650 GeV and 45 percent above 1 TeV, and 27 percent in Higgs pairs. **That is the measurement that engages M_***, and the high-luminosity run can make it.

(3) The Higgs quartic at M_*: an agreement, not a derivation. The

substrate side of the $Z^2 \to \mathrm{e}^{-1}$ identification is derived from P1–P3, and $\lambda_{\mathrm{SM}}(M_*) = b\,\mathrm{e}^{-1}$ agrees with the renormalisation-group value (the 1.4σ test of §19.7, Computation 202). That agreement is *evidential* rather than probative: the matching to the measured coupling rests on the framework premise F4, that the Standard Model is PST's effective field theory below M_*. A 1.2% coincidence across an independently fixed scale is worth something, and it is not a proof; it moves with the top mass, and exact agreement would need $m_t = 173.6$ GeV.

19.4.4. Where PST matches general relativity, on purpose

It is easy to mistake a recovery for a prediction, so the recoveries are named here. The parametrised post-Newtonian (PPN) parameters come out $\gamma = \beta = 1$ because the projected field equation is the pure Einstein equation at PPN order (§8.2), and the cosmological term's equation of state is $w = -1$ as a structural theorem (§16.5). Perihelion precession, light deflection, gravitational waves and the accelerating expansion are therefore *consistency requirements that PST meets*, not evidence that distinguishes it. The only place the theory departs from general relativity at all is near the substrate scale d_0, which is why the Casimir measurement carries the whole weight of the empirical case.

19.4.5. What the theory does not predict: the fermion masses and mixings, and why that is a theorem

The fermion mass and mixing spectrum is *not* among these results, and the reason is a theorem rather than unfinished work. The quantities are the Yukawa couplings and what follows from them: the quark and charged-lepton masses (the charged leptons in the ratio $1 : 206 : 3477$ at M_*), the neutrino masses, and the two mixing matrices, Cabibbo–Kobayashi–Maskawa (CKM) and Pontecorvo–Maki–Nakagawa–Sakata (PMNS). They sit in the contingent layer, the content carried by whichever configuration C actually sublimated, rather than in the structural layer that follows from P1–P3 alone (§14.3, §20.7).

Why this is a theorem. The Bernoulli measure on the substrate is invariant

under every permutation of its bits, and permuting the three generations is one of those permutations. So any quantity computed from the postulates alone is the same for all three generations: at leading structural order the three Yukawa eigenvalues in each sector are equal and the mixing matrices are the identity. A hierarchy breaks that symmetry, and the only thing in the theory that can break it is the particular tension distribution $T(C)$, which the postulates constrain but do not pick. The structural-scope theorem (Computation 1) states this and closes the escape routes: the mechanisms proposed for generating the hierarchy from structure, the V_7-direction constraint, the per-generation directed threshold and the structural Boltzmann weights, are each checked and each found S_3-symmetric (Computations 76, 77).

Why the distinction matters. An unpredicted quantity that a theorem places outside the theory's reach is a different object from one the theory has simply not computed yet, and PST has both kinds. The fermion spectrum is the first: no further work will derive it from P1–P3, any more than initial conditions follow from laws, and §20.7 counts what that leaves against the Standard Model's free parameters. The second kind is what the frontier of §19.9 carries, a proof of the uniform Dirac bound over its whole regime for instance (§19.9.4): not forbidden, only not yet done.

19.5. The A_F algebra identity

P1–P3, with the maximal-division-algebra (Dixon) hyperspinor envelope of §1.5 (a consequence of P1, Computation 168), imply that the substrate's spectral-triple internal algebra is $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$.

Closed in §13.7, where the chain is given in full: the parity-selected real structure delivers KO-6 (Computation 3) and is the $\mathrm{Cl}(0,6)$ chirality grading exactly (Computation 19); Furey's chain algebra delivers $\mathrm{SU}(3)_c \times \mathrm{U}(1)$ on one generation [121]; the emergent $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ supplies the $\mathbb{H}$ factor as an internal SU(2) (Computation 66); and the Dixon synthesis lands on A_F [146], with Computations 16, 18 and 19 certifying the idempotent and the invariant operators. The derivation uses Dixon's framework, not the

Chamseddine–Connes–Marcolli (CCM) classification; the two arrive at the same algebra by different routes.

19.6. Robustness of the structural-scope theorem

A natural follow-up to §14.3 is whether constraining the V_7 directional content $v : D \to V_7$ of P1 at the postulate level could override the structural-scope theorem and derive the Yukawa hierarchy from structural inputs alone. The answer is no: multiple independent symmetry arguments converge on the same scope theorem.

The simplest natural constraint, requiring $v(a)$ to lie on one of the three quaternionic sub-algebras Q_k of $\mathbb{O}$ containing $\hat{\tau}$ (the same Q_k that deliver $N_{\text{gen}} = 3$ in §14.2), preserves the S_3 symmetry of the original derivation, because the three Q_k are structurally equivalent under the G_2-stabiliser $\mathrm{SU}(3)_c$ of $\hat{\tau}$. Any $\mathrm{SU}(3)_c$-invariant residual direction therefore satisfies $p_1 = p_2 = p_3 = 1/3$ (where p_k is the squared projection of the residual direction onto $Q_k \cap \hat{\tau}^{\perp}$), giving zero Yukawa hierarchy. Breaking $\mathrm{SU}(3)_c$ at the residual-direction level to recover a hierarchy would back-propagate to break the colour gauge symmetry (catastrophic at the SM level: colour confinement is lost).

The two derivations of the structural-scope theorem, S_3 invariance in §14.3 and $\mathrm{SU}(3)_c$ invariance here, are mutually consistent and reinforcing. Postulate-level constraints on V_7 do not derive the Yukawa hierarchy. The hierarchy genuinely lives in $T(C)$, and the structural layer of P1–P3 is robust under these extensions. Supporting computations: Appendix B, rows 76–80.

19.7. The substrate-side identification of $\mathbf{Z}^2$

The substrate-side identity $S_{\text{sub}}/2^D = \mathrm{e}^{-1}$ (§12.2.2, Computation 62) is exact and independent of any SM-side identification. The observed $Z^2 \approx \mathrm{e}^{-1}$ (1.2% near-coincidence) is the SM-side measurement.

19.7.1. The CC inner-fluctuation bridge is obstructed

Direct attempts via the Chamseddine–Connes (CC) inner fluctuation of D_{sub}, targeting the substrate-Higgs Yukawa $y/\Lambda = \mu_{\text{site}}^{1/4}$ of equation (144), consistently fail. Computation 74 (single-generator) and Computation 81 (full reality term $\varepsilon' JAJ^{-1}$) give y/Λ ratios scaling as $\sqrt{2/D}$ and $\sqrt{8/D}$ respectively rather than the target D-independent $2^{-1/4}$. Computation 84 extends the family to deterministic multi-generator (requires non-integer $K = D/(2\sqrt{2})$) and Bernoulli-activated (requires $p = 1/\sqrt{2}$, not directly derivable from $\mu_{\text{site}} = 1/2$). Computation 85 (A_F-internal Higgs) reduces to the same identification problem at the moment-extraction level. Computation 86 identifies the structural cause: the Seeley–DeWitt heat-kernel asymptotic expansion that delivers the SM Higgs quartic in continuum CCM presupposes a continuous-spectrum regime; the PST substrate's discrete spectrum at the matched-scaling cutoff edge ($\pm\sqrt{D}$ all at $\pm\Lambda$) collapses this expansion entirely, leaving no Seeley–DeWitt a_4 term from which to extract λ. The standard Chamseddine–Connes (CC) machinery is therefore *structurally obstructed* for the discrete substrate triple, not merely calculationally awkward.

19.7.2. An alternative substrate-to-SM bridge

Computations 87 and 88 identify a structurally distinct substrate-side derivation of e^{-1} that bypasses the CC inner-fluctuation bridge. The asymptotic Bernoulli moment-generating function (MGF) delivers e^{-1} at the binary-distinction scale $c = -\beta_{\text{KO}} = -2$, the $\mathrm{Cl}(1,0)$ gap $2 = 1 - (-1)$, with "binary-gap-tempered" a retained label (equation (135), equation (136) of §12.3). All three ingredients are structurally derived: $\mu_{\text{site}} = 1/2$ from P1, the binary-distinction gap $2 = 1 - (-1)$ from the $\mathrm{Cl}(1,0)$ Clifford structure, and the asymptotic limit from Cramér–Bernoulli. This is a second, independent substrate-side derivation of e^{-1}: not just $S_{\text{sub}}/2^D = f(1)$ (Computation 62, spectral-action route) but also $\mathbb{E}_\mu[\exp(-\beta_{\text{KO}}\bar{X})] = \mathrm{e}^{-1}$ (Bernoulli-MGF route).

19.7.3. The substrate side of the bridge, closed

The substrate Higgs Hamiltonian $H_{\text{Higgs}}(C)$ required for the partition-function bridge is derived from three structural constraints (additivity, matched scaling and uniformity) all traceable to P1 plus matched scaling. This delivers $H_{\text{Higgs}}(C) = \bar{X}(C)$ (equation (137) of §12.3) with no free parameters, and the binary-gap-tempered substrate partition function $Z_H(\beta_{\text{KO}}) \to \mathrm{e}^{-1}$ (equation (138)). The substrate side of the bridge is closed structurally.

19.7.4. The matching, by wave-function renormalisation

The matching $\lambda_{\text{SM}}(M_*) = b \cdot Z_H(\beta_{\text{KO}})$ holds numerically (the 1.4σ test below, Computation 202). The route that reduces (B) is the wave-function-renormalisation argument of §12.5 (Computations 110–119). The matching is a field-strength renormalisation (multiplicative, not the additive Jacobian that Computation 103 excludes) and its internal renormalisation is fully derived.

- Substrate value e^{-1}.
- Field-strength selection is structural (Computation 117).
- All-orders decoupling is closed exactly by sufficiency (Computation 121), with the Morse–Bott lemma Computation 118 its shadow.
- Vertex $Z_{\Gamma 4} = 1$ (Computation 119).

Hence $Z^2 = \mathrm{e}^{-1}$, with no residual internal renormalisation. The open content is the framework premise F4 (the SM as PST's EFT below M_*) plus the empirical RGE test, with (B) thus reduced *modulo* that premise rather than parameter-free.

19.7.5. Empirical-precision test

The empirical tightness of $Z^2 \approx \mathrm{e}^{-1}$ is tested by running the SM quartic from v up to M_* and comparing $\lambda_{\text{SM}}(M_*)$ against the substrate prediction $\mathrm{e}^{-1}/4 = 0.09197$. A simplified one-loop reading (Computation 91, matching Computation 71's actual computation) gives a $\approx 5\%$ deviation, and a compact two-loop correction does not tighten it ($\approx 7\%$); these low-order readings are

truncation artefacts. The proper three-loop running with next-to-next-to-leading-order matching at the top mass and a full input-uncertainty budget, the formulas from [51] and the inputs current [233] (Computation 202), gives $\lambda_{\mathrm{SM}}(M_*) = 0.0930 \pm 0.0008$. It converges monotonically with loop order ($0.09214 \to 0.09296 \to 0.09303$, the steps shrinking) and matches $\mathrm{e}^{-1}/4$ at 1.4σ. The residual uncertainty is dominated by the top mass and its pole-mass ambiguity, *not* by RGE truncation, so the test is a converged statement rather than a low-order accident. Earlier printings said 1.0σ, anchored to a literature value of 0.0927 that the literature's own formulas do not produce; on the 2013 inputs they give 0.0924, and the top mass has fallen since. *What the test rests on beyond the running, stated because the rigidity claim does not supply it*: λ_{SM} is evaluated at $M_* = 1.285$ TeV, which is the $c = 1$ value of P3's kinetic coefficient, and $c = 1$ is assumed rather than derived (§12.2, Computation 176). Given that assumption the 1.4σ is a parameter-free prediction, which is how this book reads it; without it c would be a free parameter that the same agreement fixes, and the test would be a one-parameter fit rather than a prediction. The assumption is named in §19.9.1 for exactly this reason. On the substrate side the prediction is rigid: the field-normalisation factor is e^{-1} exactly and D-independently, and no finite-D freedom can absorb the residual (an $O(1/D)$ shift would need $D \approx 45$ against a substrate cardinality larger by dozens of orders of magnitude). The substrate-side derivation (Computation 89) is itself exact and independent of SM RGE truncation. Net: $\lambda_{\mathrm{SM}}(M_*) = \mathrm{e}^{-1}/4$ is a rigid, parameter-free prediction, consistent with the SM at 1.4σ and sharply falsifiable as the top-mass input tightens.

19.7.6. The bridge proof sketch

Computation 92 formalises the bridge identification as a seven-step structural derivation, stated once in §12.5 as the propositions F1–F9 with the Z^2 conclusion as the main theorem F11; the sketch is not repeated here. Steps (1) to (4) and (6) are the existing framework (§12.2, §1.5, A1 of §18.1, Computation 88), step (5) is Computation 89's substrate Higgs Hamiltonian, the structural "2" is the binary-distinction gap (Computation 102),

and the result $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z_H(\beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}/4$ passes the 1.4σ test of §19.7 (Computation 202). The matching itself, Bridge Premise B, is reduced to a field-strength-renormalisation identity with its substrate side closed (§19.9.5); the decoupling is closed exactly by sufficiency of the empirical mean (Computation 121), leaving the $O(1/D)$ caveat on the value.

Supporting computations: Appendix B, Computations 71–74, 81, 84–92, 100, 110–116.

19.8. Standing among substrate theories

Compared against the twenty-six other substrate-style and emergent-physics theories collected in Appendix D, PST is the only entry that simultaneously (i) posits a pre-geometric substrate (finite in internal structure, not a finite universe; the internal algebra is $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$), (ii) derives spacetime from it as a limit theorem, (iii) derives the structural ingredients of the full SU(3) × SU(2) × U(1) gauge group from the three postulates (with the full internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ via Dixon's hyperspinor synthesis), and (iv) places the three-generation count in Furey 2014's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0,6)$ as the natural reading, with the well-known *colour-versus-generation distinguishability caveat* of the octonionic-Standard-Model programme noted explicitly: the six triplets are most naturally colour states, and the three-generation reading is graded in Appendix C.4 as a structural reading, its count the three complex dimensions of $\hat{\tau}^{\perp}$, the same three as the colour count (§14.2, Computation 175), and (v) contributes the substrate-side derivation of a specific Standard-Model coupling value, $\lambda_{\mathrm{SM}}(M_*) = b \cdot \mathrm{e}^{-1} = 0.0920$ at the modal scale, matching the observed value at 1.2%, with the substrate-side factor e^{-1} derived from the foundational postulates and the SM-side matching reduced within PST to a field-strength-renormalisation identity (§19.7). Other substrate theories deliver one or two of these structural ingredients; PST delivers all five.

19.9. The remaining open frontier

PST's foundational layer (P1–P3, the substrate spectral triple, the emergence of spacetime, the gauge group, and the three-generation reading) is in place. Two items concern how PST's two headline quantitative claims meet experiment; this section states both plainly, *each as an honest bound, not a closure*. Bridge Premise (B) (§19.9.5) is reduced to a field-strength-renormalisation identity whose substrate side and internal renormalisation are fully derived (Computations 117, 119, 121). Its remaining open content is the framework premise F4 (the SM as PST's EFT below M_*) plus the empirical RGE test, not an internal gap, so $Z^2 = \mathrm{e}^{-1}$ is derived *modulo the emergence premise*, not parameter-free. F4 itself splits, and the split is worth stating because only one half is the kind of thing a proof could reach. *That no PST mode propagates below M_** is a statement about the substrate's spectrum: the sharpest objection to it is that modal sublimation is a Landau–Ginzburg transition, so Goldstone's theorem should deposit six massless modes below M_* that the Standard Model cannot house. That objection does not land, and settling it fixes a reading of P3 that had been left loose. Taking the locking as a finite-strength *potential* fails for the terms tested: with a linear or a quadratic term the angular modes come out lighter than the radial one at every strength, hence below the Higgs mass and at least a factor 10.26 short of M_*, because raising the strength raises both masses together while the ratio saturates below one. *That is not an exclusion of the branch*, and an earlier statement here said it was. The sign turns on the *curvature* of the locking term, and a quartic one reverses it above a crossover, putting the angular modes *above* M_* (Computation 170). So the branch is not ruled out, and the reading adopted below rests on P3's wording rather than on a theorem. Taking it as *direct determination*, which is what P3 says, leaves the order-parameter sector exactly one mode, the radial one, and no light modes at all. What determination does *not* remove is the *phase*, $\hat{\tau}$ being a real direction that carries none, and that phase is θ (Computation 171); it is one flat direction rather than six, and it is eaten in the electroweak mechanism rather than left in the spectrum,

so nothing massless and scalar survives below M_*. *That what remains below M_* is the Standard Model exactly* is the other half, and it is a matching to observation: it can be tested and cannot be proved. So F4 is reduced rather than closed, and stays named. The full-SM extension of the M_* boson-fermion cancellation (§19.9.6) is settled in the *negative* (Computation 120): it does not exist within PST, so M_* is irreducibly a scalar-sector prediction, a definitive *limitation* on the scope of that result, not its closure. Neither item is a claim that everything is derived; both bound what the framework delivers.

19.9.1. The structural front

No derivation remains open at the structural level. The two conditions earlier printings carried beside them, Born agreement on the law of T and non-balance, are one condition, the non-degeneracy of T's contingent content, and it is supplied by P3's typicality read as reaching that content (§1.7, Computations 193, 195), on the footing of A3; the canonical form of $T(C)$ closed with it, only the ordering of the magnitudes and $\hat{\tau}$ being load-bearing (§4.1). Both are in Appendix C. Colour-generation distinguishability is settled in the *negative* inside this theory: $\hat{\tau}^{\perp}$ is the irreducible **3** of $\mathrm{SU}(3)_c$, so no colour-invariant generation splitting exists in it at all (Computation 175), and what remains is the octonionic programme's question of where generations live if not there, which is not a derivation this book owes; its row is in Appendix C.4. The conformal identification of $\Pi_{\mu\nu}$ is closed: the conformal part of the pair sum reaches every metric-traceless tensor obeying the null energy condition (Computation 199), so no field this book contains lies outside it, and its magnitude is the state's radiation content (Computation 196); it is an identification, in Appendix C.6, and no longer a derivation this front owes. The interacting substrate vacuum has left the front: the regulated scalar sector's vacuum exists nonperturbatively, in finite and in infinite volume, by three classical theorems applied to a theory that never removes its cutoff (Computation 200), and what remains, the gauge and fermion sector below M_*, is the cutoff Standard Model's own existence problem, inherited rather than introduced; its row is in Appendix C.5. The normalisation of P3's

gradient term, the last item and an assumption rather than a derivation, is the one premise this front hands on: canonical normalisation of the emergent field, reading A of §12.2, the rival reading excluded (Computations 173, 174, 176), the Mosco limit not fixing it (Computation 177) and the derived unit of angular momentum not fixing it either (Computation 197). What Computation 201 found is that the premise fixes the coherence length itself, $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am, the gradient normalisation and d_0 being one unknown (the derivation is given once in Chapter 12). The item that waited on d_0 was waiting on itself. The premise is carried in Appendix C.8 beside F4, A1 and the Connes spectral-triple framework assumption, and nothing on this front is a derivation this book owes. The Conclusion points at this subsection rather than restating it.

None of these obstructs the foundational layer already in place; each records a closure and where it came from.

19.9.2. Localising an apparatus in the property domain

The Bell analysis of §10.5 is exact over a bipartition $D = D_A \sqcup D_B$. The passage from two *places* to such a bipartition is a corollary of the quotient construction rather than an assumption: a spacetime point is a fibre of ρ, so a region is already a set of properties, and disjoint regions give disjoint subsets of D. Two things are left over, and both are questions about apparatus rather than about geometry.

ρ is not unique, and that non-uniqueness is PST's diffeomorphism invariance, so $D_A = \rho^{-1}(U_A)$ is fixed only relative to a choice of partition. **That one is settled, and by a quantifier rather than a calculation.** PST does not need the naming to be ρ-independent, because the theorem it feeds is universally quantified over bipartitions: the tensor factorisation of §10.5 is stated for an arbitrary $D = D_A \sqcup D_B$, and no-signalling follows from it together with the Born rule, completeness $\sum_j Q_j = \mathbf{1}_B$ and cyclicity of the trace on the B factor. No step in that chain mentions ρ, a metric, a region or a distance. Every choice of ρ induces *some* bipartition, and there is no bipartition for the

theorem to fail on, so a different ρ relabels which properties constitute "the A apparatus" without touching the conclusion. Computation 151 checks the factorisation exhaustively (on all $2^{|D|} - 2$ non-trivial bipartitions at $|D| = 10$, not on a sample) and confirms that a permutation-invariant but non-product measure destroys it, so what carries the result is P1's Bernoulli measure rather than the bipartition formalism.

What is left over is the second question, and it is smaller than it looks once (30) is used rather than stepped around. Computation 156 takes it apart into three pieces, of which only one is genuinely open.

Compactness is not an extra assumption. For a spacelike pair (30) gives $|s^2| = |T - \tau_{\text{null}}|^{1/2}$, so the proper distance is $d(\rho(a), \rho(b)) = |T(\{a, b\}) - \tau_{\text{null}}|^{1/4}$, strictly monotone in the tension excess. "D_A occupies a compact region" and "D_A is tight in the tension structure" are therefore the same condition written twice, and a planted subset is recovered from the tension matrix alone, with no geometric input.

Being a device is an open condition, with a sharp margin. Write $m = \text{sep}(D_A, D \setminus D_A) - \text{diam}(D_A)$. Every perturbation of the induced distances smaller than $m/2$ preserves the property, and $m/2$ is sharp: adversarially, $0.99\, m/2$ leaves it intact and $1.01\, m/2$ breaks it. That is stability with an explicit distance to failure and it needs no measure, the same shape as Computation 152's replacement for the non-balance genericity argument. A supporting lemma makes the accounting exhaustive rather than sampled: if diam $<$ sep then every member of D_A is nearer to all of D_A than to anything outside, so D_A is exactly the $|D_A|$ nearest neighbours of each of its own members, checked against brute force over all subsets at $|D| = 14$.

The criterion needs cardinality, and then it bites. Bounded diameter with separation is on its own far too weak: in an unclustered configuration *every* nearest-neighbour pair satisfies it. Adding size fixes that. Over 200 replicas of 60 properties the accident rate runs 1.000, 0.965, 0.300, 0.040 at sizes two to five, and is 0.000 from size six upward. A device of a dozen properties is not

something an unclustered configuration produces.

What remains genuinely open is one thing: PST does not derive why T has clusters of four or more (six on the uniform null, four on the theory's own measure, Computation 212). That is a statement about the configuration rather than about the structure, so it belongs with the $T(C)$-contingent content, the same category as the Yukawa hierarchies and the Λ magnitude, and not with the open derivations. The dynamical half, why a bound state forms and persists, is inherited with the projected field theory of Chapters 12 and 13, exactly as in ordinary quantum field theory. That inheritance is forced rather than chosen, because the theory's own field cannot hold a localised object (Computation 215): in three dimensions P3's functional has no static finite-energy solution but the vacuum, every localised configuration lowering its energy by shrinking (Derrick's scaling), and on the regulated lattice the shrinking is stopped only by pinning, which needs a discreteness $(m\ell)^2$ above 47 for a single site and above 15 for a four-grain bubble where PST's is 0.0095 with the grain and 2 with the coherence length. The one long-lived localised object of the native sector is an oscillon, a breather at $1/m_h$ that lasts 10^{-23} s. An apparatus is therefore a bound state of the inherited matter sector, realised in T as a cluster that (30) reads as a device. The accident rates above were computed against a uniform null; on the theory's own measure, directions i.i.d. under the unique G_2-invariant law (Computations 141, 211), they are sharper, not weaker: every nearest-neighbour pair qualifies, a triple appears in at most a few percent of replicas, and nothing of four or more ever does (Computation 212). The same computation says what a single point carries: a fibre of ρ is a clique of pairwise-null properties, hence a set of directions with one pairwise inner product, and such a set in V_7 has at most seven members (fewer for an obtuse null angle), while on the theory's measure an exact coincidence has probability zero, so fibres are singletons almost surely. A device of a dozen properties is therefore a region, never a point, and a configuration-level object on the theory's own typicality. The margin result is measure-free and is the part to lean on. The Bell conclusion depends on none

of it.

19.9.3. The black-hole interior

The strong-field limit (§8.3) closes two questions and leaves three open, and it is worth separating them because the open ones are routinely assumed to carry the closed ones. What is closed is the horizon's relation to light: both the light cone and the horizon are the same null threshold τ_{null}, so a photon is not held back by a force it has no mass to feel. That much follows from (30) and needs nothing further. What is also closed, and was carried here as open until the projected action's validity condition was evaluated rather than assumed (§8.3.3), is the thermodynamic entropy.

What is open is the interior, one of the two things that go by the name of black-hole entropy, and the information question, which nothing here touches. The sentence "no such count exists here" is true of that one and was never true of the other.

The thermodynamic entropy follows, and no part of the argument is a count. Wald's Noether-charge construction [207] assigns an entropy to a stationary black hole in *any* theory of gravity derived from a diffeomorphism-invariant Lagrangian, as a local geometrical integral over the horizon, with no microstate count anywhere in the argument; for the Einstein–Hilbert Lagrangian it returns $S = A/4G$. PST derives exactly that Lagrangian, with exactly that coefficient, at equation (50). What used to be missing was not thermodynamic machinery but something prior, whether the projected field equations reach a horizon at all, and §8.3.3 settles it: at $r = r_s$ the expansion parameter $(k\, d_0)^2$ is 1.9×10^{-42} for a solar mass ((66), Computation 216). In that region the projected equations *are* Einstein's, so every general-relativistic black-hole exterior solves them; the bifurcation surface sits at $r = r_s$, inside the covered region; and the hypersurface Wald's first law is integrated over lies wholly in the exterior, never entering the region the expansion gives up on. The theorem applies, and returns $A/4G$.

The truncation costs nothing here, which is worth saying because it is not

obvious. The corrections named at (49) are $\mathcal{O}(R_{\mu\nu}R^{\mu\nu},\, R^2)$, and every term of both their field equations and their Noether-charge integrands carries a factor of R or of $R_{\mu\nu}$, which vanish identically on the solution. So the black hole solves the *truncated* equations exactly, and its Wald entropy is $A/4G$ with no $(d_0/r_s)^2$ correction at all. The one invariant not named there, $R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$, is in four dimensions the Gauss–Bonnet combination up to the two just eliminated, and its Noether charge works out to $64\pi^2\lambda$, independent of the mass: an additive constant rather than an area term, and of order unity against $S/k_B = 1.0 \times 10^{77}$ for a solar mass. PST's own cosmological constant spoils exact Ricci-flatness and costs a further relative 10^{-119}. All four statements are checked in Computation 154, against the same Noether-charge formula calibrated to return $A/4G$ from (50).

Three qualifications, so the result is not oversold. It *reproduces* $A/4G$ rather than predicting it: G enters PST as a consistency relation among (m_h, M_P, v) rather than as a first-principles output (§8.5), so the Planck mass is an input on that route. Noether-charge entropy is a property of the classical Lagrangian; it says nothing about what the substrate is counting, which is the next problem and is untouched by any of this. And the idealisation is the standard one: the theorem wants a stationary solution with a bifurcate Killing horizon, which a hole formed by collapse has only as the late-time limit of its exterior. PST inherits that idealisation from general relativity rather than adding to it, and inherits unchanged the question of whether $r = r_s$ is an event horizon rather than merely an apparent one, since that turns on the interior being censored, and the interior is the one place the projected equations do not reach.

The statistical entropy is a research programme, and it is in tension with the theory's sharpest prediction. A substrate account would have to count near-threshold configurations against horizon area, and no such count exists here. Three obstructions stand in front of it, and they are structural rather than technical. P1's measure is a product over D, so substrate entropy is extensive in $|D|$ and therefore in four-volume, not in area. The projection kernel $K(x, C') = \delta(x - \rho(C'))$ is supported on the bulk, and no construction

in this book puts substrate data on a surface. And the lengths PST owns are the coherence length $d_0 = 2.2$ am and the grain $\ell_* = \hbar c/M_* \approx 1.5 \times 10^{-19}$ m, neither of them near ℓ_P; an area law quantised in either would miss $A/4\ell_P^2$ by tens of orders of magnitude. That last point is not a gap to be repaired but a commitment the theory has already made in the other direction: $d_0 \gg \ell_P$ is what fixes the Casimir correction's size, a length seventeen orders above ℓ_P and an amplitude thirty-four orders above a Planck-scale one (§15.3.1), so a d_0-quantised horizon count would trade the theory's one derived laboratory number for its entropy. A separate point concerns the dissolution of the *initial* singularity, which Table 2 claims on the grounds that loci are not primitive, is an argument about the beginning of the instantiated manifold; for a black hole's central singularity it is made in a different form by Computation 216: a point carries at most seven properties, so the collapsed content occupies a region whose minimal diameter the packing bounds give in local units, and the singularity as a point has no substrate counterpart. What that region does, the theory does not say; the projected equations stop at $r_{\text{break}} = 37$ pm for a solar mass, a region of 10^{25} grains. None of this obstructs the gravitational sector already derived; all of it marks where a substrate account of black-hole thermodynamics would have to begin.

19.9.4. Beyond the two quantitative bounds

One analytic residual remains open, and it is narrower than it was. The equicoercivity A7 is a theorem given A1, A2, A4 and A3: its equidistribution condition A3, the closure-rate inequality $\eta_{K,R} < d_0^2/4$ on the emergence map ρ_n (Computation 129), is supplied by the maximum-entropy sublimation of P3 rather than adopted as a premise (§18.1.6). The Dirac-convergence step is closed at the relaxed-criterion level, and the surviving open item is the one remaining quantitative refinement: a proof, rather than a derivation at the maximiser and a scan elsewhere, that the constant of its uniform bound on $k \leq k_D$ is what the computations find. The existence of a finite bound there is an analytical theorem, and its mechanism is now identified: the deficit is governed by $m(k)/N$, which the matched scaling drives to zero at rate $D^{-1/4}$

(equation (282), Computation 150). The sharp value, 3.6736 at its corner, is now an explicit algebraic number, and the problem it sits in is no longer opaque (§19.2.1, Computation 157): $[J_z, T] = 0$ removes the Dirac operator, leaving a maximum over ratios of factorials on an integer lattice; the case $k = 1$ closes exactly at ratio$^2 = 4(N-1)/N$, so $C(1) = 1$; and the constant has been rechecked to $k = 80$ rather than the $k = 20$ Computation 150 reached, with the maximum still 3.6736 at $(k, N) = (6, 12)$, now certified in exact arithmetic over $k \leq 20$ (Computation 203). The discrete maximum over a, which generates the sawtooth, is in closed form for every m (Computation 203, equation (276)): the single-monomial problem is closed, its $1/N$ deficit cancelling for $m \geq 2$ by a proved expansion whose N^{-2} coefficient is explicit (Computation 204, equation (277)). The two-monomial bridge's deficit, which rose slowly with N to 15360 with no plateau in sight (Computation 206), is a limit after all, in closed form and approached from below as $N^{-1/3}$ through the Airy edge of the bridge's own chain (Computation 213, equation (278)): below the derived 3.6736 for every fixed $k \leq 482$, so the closure rate carries no logarithm; and the finite-N sharp value is an explicit algebraic number at its corner, $(6, 12)$, which is the maximum because it is the one point where a headroom of four meets a balanced bridge (Computation 214, equation (275)). What remains is that the limit grows slowly with k, so the constant is not uniform in k without end, and that the supremum is derived at its maximiser and scanned elsewhere rather than proved over the whole regime. The open items of the foundational and dynamical layers are the ones listed in this section and in §19.9.1; nothing is carried elsewhere. The closures and contingencies already established elsewhere in the book:

- The KO mod 8 $\to \beta_{\mathrm{KO}}$ identification is resolved by Computation 102. The load-bearing "2" is the spectral range of a single binary distinction, $2 = 1 - (-1)$, forced by P1's binary alphabet, equivalently the one-bit Clifford $\mathrm{Cl}(1, 0)$ range. The KO mod $8 = 2$ value is flagged as a structurally independent coincidence; see §12.3.
- The full measurement postulate / Born rule + collapse is in §10.3, resting

on the von Neumann coupling and the assumed unitary propagation rather than on Gleason alone; the fermion sector's two residuals, the slice-wise CAR limit and microcausality W4, are closed in §11.3 given A1 to A4 with A3 supplied by P3, with only textbook axiomatic-field-theory work remaining. A7's supplied condition A3 (stated above; §18.1.6) requires that the metric-covariant unevenness of the realised cloud stay within the closure margin $\eta_{K,R} < d_0^2/4$. It is secured in both channels. The fluctuation channel is settled unconditionally by McDiarmid (Computation 125). The rate that P3 supplies for the emergence-determined ρ_n is corroborated on the i.i.d. surrogate (Computations 70, 83); it is a deterministic equidistribution statement rather than a mixing rate (Computation 129), and not a bulk-density freedom, which is a coordinate artefact (Computation 134). This is a typicality (maximum-entropy) selection internal to P3, an analogue of maximum-entropy or generic-initial-data selection, not a fourth postulate: sublimation instantiates a μ-typical configuration, whose cloud is equidistributed at the rate A3 states and whose content lies off every closed set of positive codimension.

- **Intermediate coupling.** Whether a weak coupling leaves an interference signature is closed, in the negative (Computation 144): a crossing that selected would violate the Englert–Greenberger–Yasin relation [203, 204] downward, which experiment excludes [205], so crossings are redundant with decoherence and PST reproduces standard quantum mechanics throughout this regime. What remains open is stated in §10.3.2.5: the coupling at which the einselected pointer basis, rather than (98), is doing the separating.

- Three quantities are $T(C)$-contingent rather than open derivations (§14.3, Computation 82): the Λ_{obs} magnitude and its $\sim 10^{-120}$ hierarchy, the Yukawa hierarchies, and Cabibbo–Kobayashi–Maskawa (CKM) mixing; the d_0 magnitude is not among them, being fixed by the normalisation premise (Computation 201). The cosmological-constant *sign* and *equation of state* $w = -1$ remain structural theorems (§16.5).

- The $v \to M_* \to v$ running self-consistency that lets the two electroweak

observables (m_h, v) collapse to one (with v following from m_h) is tested rather than assumed since Computation 202 part 9 (§12.2): running $\mathrm{e}^{-1}/4$ back down from M_* returns $v = 247.1$ GeV from m_h, 0.4 percent high, the deficit at M_* seen from the other end. What remains assumed is that the run-up to M_* and back is insensitive to the v being solved for; the 0.4 percent bounds that sensitivity without proving it absent. The stakes are low: $Z^2 = \mathrm{e}^{-1}$ pins $\lambda_{\mathrm{SM}}(M_*)$ without using v as an input, so the headline agreement is insensitive to this self-consistency.

- The Casimir correction is falsifiable through its *exponent* and, now that the normalisation fixes d_0, through its *amplitude*, but not through non-observation (Chapter 15): the derived $d_0 = 2.2$ am puts the signal at 10^{-19} of the force, so a null result at any reachable precision is what the theory expects and merely lowers the $d_0 \lesssim 2.3$ nm bound. The falsifiers are a measured short-distance power law inconsistent with the parameter-free d^{-6} exponent, and a d^{-6} component resolved at any reachable amplitude, which would refute the canonical normalisation and with it M_*.
- The primordial perturbation spectrum's amplitude and tilt are not derived. The Bernoulli measure gives fluctuations that are Gaussian at leading order, with the skewness falling as $(d_0/\lambda)^{3/2}$ and the excess kurtosis as $(d_0/\lambda)^3$, and a spectrum that is white at the grain (Computation 217); the Harrison–Zeldovich power and the measured tilt $n_s - 1 = -0.035$ [71] belong to the transfer through the near-threshold epoch (§16.5), which this book does not carry out.

19.9.5. Bridge Premise (B), reduced

The substrate-side factor $Z_H(\beta_{\mathrm{KO}}) \to e^{-1}$ is derived from P1–P3 (Computation 100), and its identification with the SM-side ratio $\lambda_{\mathrm{SM}}(M_*)/b$ is *reduced* by the four-step argument of §12.5, given there once: (B) is a wave-function renormalisation rather than a coupling-matching Jacobian (Computations 101, 103, 110); P1's product structure forces its kernel into the substrate-factor form, the Gaussian excluded (Computation 111), with $\beta_{\mathrm{KO}} = 2$ the binary-

distinction gap (Computation 112); its substrate-side value is the embedding norm of the symmetric Higgs vacuum, the standard-EFT two-point reading excluded because the law of large numbers sends it to zero (Computation 117); and the total reduction is a theorem, $\bar{X}$ being the unique soft mode (Theorem F10; Computations 115, 116), with the vertex renormalisation trivial to all orders (Computation 119) and the decoupling closed exactly by sufficiency of the empirical mean (Computation 121). What is not derived is the SM-side content: $\lambda_{\rm SM}(M_*) = b\,Z^2$ is Wilsonian threshold matching resting on F4, tested by the three-loop running (the 1.4σ test of §19.7, Computation 202), converged in loop order and with no substrate-side knob to absorb a miss. So (B) is reduced, not closed: the open content is F4 plus the empirical test, not an internal renormalisation step. The heat-kernel CC route is obstructed for the discrete triple (Computations 85–86); the construction is internal to PST, not inherited.

19.9.6. The full-SM extension, settled in the negative

The scalar-sector self-mass relation $m_h^2 = (3/2)M_*^2/(16\pi^2)$ is a parameter-free one-loop result for the Higgs self-loop in isolation (§12.2, Computations 93, 97). Its extension to the full Standard Model, a cancellation of the EFT-layer top, W and Z quadratic sensitivities that would let $M_* = 4\pi m_h\sqrt{2/3}$ stand as a full-theory prediction, does *not* exist within PST (Computation 120). The three obstructions are set out once in Chapter 12: the substrate cancellation acts in a layer disjoint from the SM loops (Computation 105), its sector-blindness not transferring it (Computation 107); the Veltman bracket has no zero crossing at M_*; and the bosonic top-partner sector a cancellation would need is excluded by the LHC, by the "no new light states" premise of §15.11 and by F4. Hence M_* is irreducibly a *scalar-sector* prediction: a definitive limitation on its scope, not a closure of it, which removes the item from the open register with M_* finally scoped.

19.10. Summary

This chapter has catalogued the research problems that PST resolves and the present status of each. Eight results are established at the structural level, each with its conditions and any open quantitative refinement stated explicitly: the foundational object and the emergent Lorentzian geometry; gravity and cosmology; quantum mechanics, including the Born rule and the measurement postulate; the Standard Model gauge group, with its anomaly-cancellation and hypercharge consistency; the Dirac-convergence step; the octonionic SU(3) structure inside $\mathrm{Cl}(0,6)$; the matter content; and the A_F algebra identity (Conjecture 1). Among substrate theories PST is the only entry that simultaneously derives spacetime, the gauge group, and a falsifiable scalar-sector prediction from a fixed set of postulates, and places the three-generation count as the natural structural reading (§14.2; its former condition, the Dixon block structure of Computation 106, is withdrawn, Appendix C.4). Two items carried the quantitative weight: the matching premise of the $Z^2 = \mathrm{e}^{-1}$ bridge (§19.7, open), and the full-Standard-Model scope of M_* (settled in the negative, a bound rather than a closure, Computation 120). Both are reduced to sharp, named outcomes rather than left diffuse, and both are collected in §19.9. Beside them the frontier carries five named premises (canonical normalisation, F4, A1, the spectral-triple framework, the odd-$|D|$ parity selection) and three residues (localising an apparatus, the black-hole interior, the uniform bound beyond $k = 482$), each stated once there and indexed in Appendix C.8.

20. Conclusion

A short account of what the book claimed, what it established, and what it left open. The last of those is the most important, and it is stated plainly rather than buried.

20.1. 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.

20.2. 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, together with the two further postulates P2 and P3 and one extremal selection principle (maximality over the Hurwitz classification, applied to the direction algebra and to the hyperspinor envelope) that is a consequence of P1 rather than a commitment beside it (Computation 168), 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, §4.12). Each downstream concept follows from the preceding ones; the chain introduces no free numerical parameters, only the single named structural selection above, flagged where it is used. 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, or the specific matter content of this universe down to the Yukawa hierarchy and Cabibbo–Kobayashi–Maskawa (CKM) matrix. These depend on the particular tension configuration C that produced this universe and require knowledge of C beyond its structural form (§14.3, structural-scope theorem). What PST derives structurally is the full $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ gauge group and one generation of Standard-Model fermion content; the generation count $N_{\mathrm{gen}} = 3$ follows as the natural structural reading of the directional content of V_7 (§14.2) rather than entering as a free parameter.

20.2.1. Standing among substrate theories

Of the twenty-seven substrate-style and emergent-physics theories surveyed, PST is the only one that meets all five of the axes on which they are compared; the others meet one or two. The five axes are stated and argued in §19.8, and the full comparison is tabulated in Appendix D.

20.2.2. The Dirac-convergence constant

The Dirac-convergence step is established at the relaxed-criterion level that suffices for the Standard-Model derivation chain (Chapter 19). One quantitative refinement remains: a proof over the whole operative regime of the uniform

bound on $k \leq k_D$ (§18.3), whose existence and $D^{-1/4}$ mechanism are analytic, whose large-N limit is in closed form (Computation 213) and whose sharp value is derived at its maximiser and scanned elsewhere (Computation 214).

20.3. What the emergence chain establishes

The core argument of the book 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 modal reorganisation by which the supra-threshold configuration is instantiated as geometry, represented mathematically as the functor $\Phi : \mathbf{S} \to \mathbf{G}$.

The functor Φ produces a Lorentzian manifold (M, g). The signature $(+, -, -, -)$ follows by hyperbolic well-posedness of the Goldstone wave equation (Computation 2, §4.12): Lorentzian is the unique signature admitting a well-posed Cauchy problem and a conserved Noether current. The Dirac-convergence step to the unique spacetime Dirac operator is established at the relaxed criterion (§18.3; Chapter 19). 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's 2-plane projection, the circle $\mathcal{V}_{\hat{\eta}} \cong S^1$, 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 (28). The tension-to-metric map is equation (30). The stress-energy tensor is equation (39). The conservation law is equation (48). The Noether charge is equation (80). The chain is not a qualitative narrative but a sequence of mathematical constructions, each of which follows from the previous given the three postulates of §1.7.

20.4. 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, 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 $\mathcal{V} \cong S^1$. What PST fixes structurally is the location of the vacuum, the sign, and the equation of state $w = -1$ of the observed accelerating component, all carried by the ongoing-sublimation density (§16.5). The *magnitude* of the static valley depth renormalises like any vacuum energy and is not derived (§9.1). So the 120-orders mismatch is relocated and disarmed rather than numerically resolved.

Why do orbiting bodies orbit? Because the instantiated vacuum projects onto the 2-plane as 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.

20.5. 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 quantise 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.

20.6. 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 (28); the functor Φ is constructed in three explicit steps via the realisation map ρ and the tension-to-metric assignment of equation (30); the projection operator Π is equation (31); the field equations follow as equation (48); the vacuum manifold radius is $r_0 = \sqrt{\varepsilon/(2b)}$; the Noether charge is equation (80). Each of these is a mathematical statement, not a metaphor.

The primary quantitative prediction of the theory is a correction to the Casimir effect that falls as d^{-6} against the standard d^{-4} (equation (208)), with its coefficient $\xi = 90/\pi^2$ parameter-free (Computation 73) and its amplitude fixed by the coherence length the canonical normalisation puts at $d_0 = 2.2$ am (Computation 201): 10^{-19} of the force at 20 nm, nine orders below the 2.3 nm the data bound, so a null result at every reachable precision is what the theory expects, and a d^{-6} component resolved at any reachable amplitude would refute the normalisation premise. The numbers, the ratio test and the sensitivity at the bound are in §19.4.3. 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. So the dimension PST predicts *decreases*, from high values at short distances to four at large scales. Earlier printings called this qualitatively consistent with Causal Dynamical Triangulations (CDT) [31]; it is not, and the claim is withdrawn. CDT reports the opposite flow, a spectral dimension near two at the Planck scale rising to four at large scales, and the same dimensional *reduction* appears in asymptotic safety and in most of the other approaches Carlip reviews [231]. The one published result with PST's direction is in causal set theory: diffusion on the full sprinkled causal set gives a spectral dimension that *rises* at short scales [229], which Carlip's review records as the outlier; the same group later traced the rise to the unbounded number of nearest neighbours a Lorentz-invariant discretisation has, and found dimensional reduction again once the diffusion is confined to a spatial hypersurface [230]. PST's prediction therefore stands on the minority side, with the full-causal-set result and against CDT. Whether

PST's near-threshold dimensionality is also a spectral dimension in Carlip's sense, the return probability of a diffusion on the projected geometry, is a computation, and Computation 207 performs it exactly on the i.i.d. surrogate of Chapter 18 (Computations 70, 83): the Boolean Dirichlet form's walk, one flip per unit of Mosco time, pushed to a flat Euclidean bulk that stands in for M by Gaussian displacements with the A4 covariance that stand in for the emergence-determined realisation. The answer depends on what the probe resolves, and the dependence is the whole content. Above the coherence length d_0 the spectral dimension is four, on a window that grows in proportion to $|D|$: the Mosco limit read as a diffusion exponent. At the resolution d_0 itself, the finest at which ρ distinguishes points, it never exceeds 4.5, an overshoot smaller than the cubic lattice's, and falls to zero below d_0 as on any discrete structure. At the grain $\ell_* \approx d_0/14.5$ it rises above four, to about 23 at a few coherence lengths, before settling; the height is logarithmic in d_0 over the resolution and unbounded for exact return on the hypercube. The mechanism is the one the causal set group identified: a walker's first step can go in any of $|D|$ directions, so the return probability falls exponentially until the diffusion has spread past the resolution, and resolving at d_0 bounds the effective valency as confining the causal set diffusion to a hypersurface does. So in the sub-d_0 regime the withdrawn sentence was about, the spectral dimension and the direction count flow the same way, both falling to four as the scale grows, and PST stands with the full-causal-set result and against CDT as a computed statement rather than an asserted one; reduction to two appears nowhere in these models, and the direction count and the diffusion exponent remain different quantities. The statement carries one condition, and Computation 210 supplies it on the book's own finite model rather than the surrogate: on Computation 9's bridge, where the binomial multiplicity of the Walsh ladder is absorbed into the measure, the spectral dimension is three between the truncation and the compact scale, falls to zero below the truncation, and never rises; keeping the binomial multiplicity restores the rise. So the rise between the grain and d_0 is carried by the Walsh modes the bridge reweights away, and a probe sees it exactly when it couples to the substrate's

walk rather than to the emergent geometry, which is the resolution statement above in the bridge's terms.

20.7. The scope and limits of the theory

20.7.1. What is foundational, and what is inherited

PST has three foundational postulates: P1 (property differentiation), P2 (asymmetric vector tension), and P3 (modal sublimation past a Landau–Ginzburg threshold). From these, the substrate spectral triple, the internal algebra A_F, the Standard Model (SM) gauge group, and the substrate-side matched-scaling coefficient e^{-1} (Computation 100) are all derived; the Mosco limit to spacetime $M = \mathbb{R} \times S^3$ closes under A1–A7 rather than under P1–P3 alone, its equidistribution condition A3 supplied by P3's maximum-entropy sublimation and its matched-scaling condition A1 carried as architectural input and corroborated on the i.i.d. surrogate for the emergence-determined ρ_n (§18.1). The three-generation count $N_{\mathrm{gen}} = 3$ is a structural reading graded in Appendix C.4 (§14.2), whose stated condition, the Dixon block structure of Computation 106, is withdrawn; it is not a consequence of P1–P3 alone.

For the SM-side matching at the modal scale M_*, PST invokes a partition-function-level analogue of the Chamseddine–Connes spectral-action correspondence [49, 50] (a construction novel to this book, since the heat-kernel Chamseddine–Connes [CC] machinery is structurally obstructed for the discrete substrate triple). Bridge premise (B) is reduced within PST to a field-strength-renormalisation identity with its internal renormalisation complete, so $Z^2 = \mathrm{e}^{-1}$ is derived modulo the emergence premise plus the renormalisation-group equation (RGE) test at 1.2%, 1.4σ: parameter-free, but not premise-free (§19.7).

The structurally-derived content of PST from postulates P1–P3 is large and concrete:

- the emergence of spacetime $M = \mathbb{R} \times S^3$ as the Mosco limit of the substrate's Boolean Dirichlet form, a theorem under A1–A7 whose equidistribution

condition A3 is supplied by P3's maximum-entropy sublimation, and whose matched-scaling condition A1 is architectural input (§18.1);

- the Standard-Model gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ as the unimodular unitaries of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ (Chapter 13, Computation 6);
- the internal algebra A_F itself, derived from the substrate's Cl(0,6) chain-algebra structure (§13.7, Computations 3, 19);
- the Higgs mechanism with the modal-condensate identification $\phi_{\text{Higgs}} = \Pi(\psi)$ and $m_h^2 = 4\varepsilon$ (§15.2);
- Einstein's field equations as the leading Seeley–DeWitt term of the spectral action (§8.5, Computation 5);
- the Born rule from Gleason's theorem applied to $L^2(\mathcal{P}(D), \mu)$ on the Bernoulli measure (§10.2);
- the existence of fermionic excitations via the Boolean–canonical-anticommutation-relations (CAR) isomorphism on the distinction primitive (§11.2);
- the cosmological principle as a substrate theorem from permutation invariance of μ (§16.2);
- Geroch time-orientability of the emergent manifold from P3-directed threshold plus permutation invariance plus Mosco convergence (Chapter 5);
- Lorentzian signature $(+,-,-,-)$ from hyperbolic well-posedness of the Goldstone wave equation (Computation 2);
- the four-dimensional spacetime $n = 4$ from the minimal-Connes-spectral-triple split given the real K-theory (KO) dimension-six substrate, which is enforced by the Connes spectral-triple framework assumption rather than by P1–P3 alone;
- the d^{-6} Casimir-correction power law (with coefficient $\xi = 90/\pi^2$ derived

under the Gaussian projection kernel, the one modelling choice C.2 records, the exponent itself kernel-independent, Computation 73);

- the cosmological-constant sign and equation of state $w = -1$ from the steady-state non-dilution mechanism (§16.5);
- the substrate-side factor $Z_H(\beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}$ of the field-normalisation identity $Z^2 = \lambda_{\mathrm{SM}}(M_*)/b(M_*) \approx \mathrm{e}^{-1}$, derived structurally from P1–P3 (Computation 100). Bridge premise (B) reduces within PST to a field-strength-renormalisation identity resting on premise F4 (the SM as PST's effective field theory [EFT] below M_*; §12.6.2.1) plus the RGE test rather than a free parameter, so that $\lambda_{\mathrm{SM}}(M_*) = b \cdot \mathrm{e}^{-1} = 0.0920$ matches the observed value at 1.2% (full account, §19.7).

This is derivation *from three foundational postulates* (P1, P2, P3) of the structural form of physical reality: spacetime, gravity, the gauge group, the Higgs and fermionic matter. The three-generation count is not on that list: it is a structural reading whose stated condition, the Dixon block structure of Computation 106, is withdrawn, and it is graded in C.4, not C.3. One extremal selection principle is used along the way, and it is not a fourth commitment: Computation 148 reduced the three once entertained to one and Computation 168 derives that one from P1, since a proper sub-object costs a selector P1 does not supply (§1.7); the spacetime chain's own conditions, A1 to A7 with A3 supplied by P3, are the ones stated at the head of this chapter.

For the SM-side matching at M_*, PST extends the Chamseddine–Connes spectral-action programme [49, 50] in two ways. It provides a precausal foundation beneath the spectral triple, since P1–P3 deliver the substrate spectral triple itself (Computations 3, 19). And it contributes the substrate-side coefficient e^{-1} as a genuinely new derivation (Computation 100), with bridge premise (B) reduced within PST so that $Z^2 = \mathrm{e}^{-1}$ is derived modulo the emergence premise plus the empirical RGE test: parameter-free, but not premise-free (§19.7).

Intellectual honesty also requires stating clearly where each derivation stands and what remains contingent. The emergence chain reaches angular momentum and the topology of the vacuum as step 8. The gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ arises as the unimodular unitaries of the internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ (§13.7); the three-generation count is delivered by the V_7-directional mechanism of §14.2. The Born rule is derived from the Bernoulli measure and Gleason's theorem in §10.2. The full measurement postulate (the collapse rule) and the fermionic-sector treatment both hold at proof-detail level. The measurement postulate rests on three closures and one condition of §10.3: an explicit von Neumann measurement Hamiltonian, decoherence inheritance, threshold-crossing uniqueness with the irreversibility theorem, and agreement with the Born rule, which holds where the einselected pointer basis rather than (98) separates the outcomes (§10.3.2.5). The fermion sector rests on the three closures of §11.3, two of them given A1 to A4: CAR-to-continuum Mosco convergence, Wightman axioms verification, and the substrate-derived Dirac equation. The electroweak modal scale is fixed, on the canonical normalisation, as $M_* = 4\pi m_h \sqrt{2/3} \approx 1.285$ TeV via a parameter-free one-loop derivation of the scalar sector; the constant-function substrate vacuum mode is structurally distinct from the emergent Higgs doublet in the projection chain (Computations 93, 97). The coherence length d_0 is not pinned by the postulates alone; it is fixed by the same canonical normalisation on which M_* rests, at $\sqrt{2}\,\hbar c/m_h = 2.2$ am, so the Casimir amplitude is derived under that one premise (§15.3, Computation 201). The dimensionless coefficient ξ is derived from the substrate's binomial-to-Gaussian Boolean kernel in Computation 73, with $\xi_{\mathrm{substrate}}(|D|) = (90/\pi^2)(1 - 2/(3|D|))$ converging to the Gaussian value 9.12 at the matched scaling. The canonical measure μ on configuration space was derived from first principles in Chapter 2, §2.5. The gradient structure on $\mathcal{C}$ required by the functional calculus is delivered by the Boolean Dirichlet form E^B on $L^2(\mathcal{P}(D), \mu)$, together with the Mosco-convergence theorem of §18.1. The topology needed to define $\nabla_{\mathcal{C}}$ is the symmetric-difference topology on $\mathcal{P}(D)$ induced by P1's distinguishability primitive.

20.7.2. Structurally contingent versus derived

Every physical theory has some contingent inputs. The Standard Model takes 19+ free parameters as empirical inputs with no structural derivation: the three gauge couplings g_1, g_2, g_s; the nine quark and charged-lepton Yukawa couplings; the four CKM parameters (three angles plus one phase); neutrino-sector parameters (masses plus Pontecorvo–Maki–Nakagawa–Sakata [PMNS]); the quantum chromodynamics (QCD) vacuum angle θ_{QCD}; the Higgs mass and vacuum expectation value; and the cosmological constant Λ. PST, by contrast, identifies a small number of items that genuinely live in the configuration content $T(C)$ rather than the structural layer of P1–P3. The structural-scope theorem of §14.3 pinpoints these as:

- the Yukawa mass hierarchy of the quarks and charged leptons,
- the CKM mixing matrix,
- the cosmological-constant magnitude Λ_{obs} (sign and equation of state $w = -1$ are structural theorems of §16.5; the specific magnitude $\sim 10^{-120}$ below Λ_{P} is configuration content via the $T(C)$ functional, Computation 82),

together with the analogous neutrino-sector content (PMNS, neutrino masses). Four families of contingent items rather than nineteen-plus; the coherence length d_0 is not among them, being fixed by the normalisation premise instead (Computation 201). Everything else the SM treats as a free parameter, PST derives structurally from P1–P3: the gauge group $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$, the Higgs mass relation $m_h^2 = 4\varepsilon$, the cosmological-constant sign and equation of state $w = -1$, and so on.

The remaining contingency is not a defect. Foundational theories are expected to distinguish structural derivation (what is necessary) from configuration content (what is contingent), in the same way that General Relativity does not derive the universe's initial conditions and Quantum Mechanics does not derive a particle's specific wave function. The honest statement that a small set of items lives in $T(C)$ is more rigorous than claims to derive "everything from nothing" that, on closer inspection, deliver a landscape with no specific

predictions.

These are contingencies, not weaknesses; they delimit the configuration content that any instantiated reality must supply. A theory that identifies the correct primitive and constructs the correct emergence chain from it will inevitably leave such configuration content undetermined by structure alone, which is the nature of a foundational theory.

20.8. Open questions and the next research

Every open item is stated once, with its conditions, in §19.9 and §19.9.1, and in the register of Appendix C.8, and none of it is restated here. Its shape is this: two bounds on the headline quantitative claims, one reduced and one settled in the negative; four named premises carried rather than derived (canonical normalisation of the emergent field, which is the one that fixes a laboratory number; F4; A1; the Connes spectral-triple framework; the odd-$|D|$ parity selection behind KO-6); nothing open at the structural level; and the three residues the frontier's subsections carry, which clusters $T(C)$ has, what the 37 pm region at a hole's centre does (with the statistical entropy and the information question), and a proof of the sharp constant over the whole regime beside the limit's slow growth in k. What follows is what comes next rather than what is missing.

The experimental front. §19.4.3 names the three places the theory meets measurement. The Casimir d^{-6} correction is exact in form and its amplitude is derived, and that amplitude is far below present reach; what a measurement at $d \approx 50$ nm at the 1% level can do is lower the diagnostic bound on d_0, or, by resolving a d^{-6} component at any reachable amplitude, refute the normalisation premise and with it M_*. The live confrontations are the other two: the Higgs transverse-momentum spectrum, where the modal scale's kernel on the doublet predicts a deficit of 23 percent above 650 GeV (Computation 205), and the 1.4σ agreement of the Higgs quartic at M_* on current inputs, which the top mass more than any other input can tighten or break. Beyond these, the

contingent layer $T(C)$, Yukawa hierarchies, CKM and PMNS mixing, the cosmological-constant magnitude, is by construction not pinned by P1–P3 alone, and a derivation of any of it would require structure beyond the three postulates.

20.9. 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.

20.10. 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 condition is derived from these three via Mosco convergence of Boolean Dirichlet forms. Most structural form beyond those three commitments follows as a consequence, under one extremal selection principle that is itself a consequence of the first postulate rather than a commitment beside it; what does not is carried explicitly, and §19.9.1 and §19.9 name it: the bridge premise (B), reduced rather than closed, and four named premises: the framework premise F4, the matched-scaling condition A1, the Connes spectral-triple framework assumption, and the canonical normalisation of the emergent field on which M_* and Z^2 jointly depend. Configuration content remains contingent (the Yukawa hierarchies, the mixing matrices, and the magnitude of the cosmological constant) and no initial conditions are imposed. What matters at this stage is that the foundational structure is correct: the primitive is genuinely irreducible, the chain of entailments is logically tight, and 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. The theory makes concrete experimental predictions that distinguish it from all current frameworks: a d^{-6} Casimir correction, exact in form and far below present reach, and a deficit in the Higgs transverse-momentum spectrum above 650 GeV. 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.

Appendix A: Notation Conventions

The following symbols recur throughout the book. The "Where" column points to the place each is first introduced, a section such as §4.6, or a whole chapter where the symbol is the chapter's subject.

A.1. Substrate and primitives

Symbol	Meaning	Where
(D, δ)	Discrete distinction substrate: a set D with the distinction primitive δ	Ch. 2
$\lvert D\rvert$	Substrate cardinality (number of distinction bits), generic letter D also used by abuse for $\lvert D\rvert$ in dimensional analysis	Ch. 2
$\mathcal{P}(D)$	Boolean configuration space of subsets of D	Ch. 2
$a \in D$	Single distinction (Boolean bit)	Ch. 2
C	A precausal configuration, $C \subseteq D$	Ch. 2
$T(C)$	Asymmetric-tension functional on $\mathcal{P}(D)$ (distinct from the Dixon algebra $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$, Ch. 13, and the oblique parameter T)	Ch. 2
μ	Bernoulli product measure on $\mathcal{P}(D)$, $\bigotimes_a \mathrm{Bern}(1/2)$	Ch. 2
V_7	7-dimensional algebra carrying the substrate's directional content, identified with $\mathrm{Im}(\mathbb{O}) \cong \mathbb{R}^7$ via Hurwitz's classification	Ch. 2, Ch. 13
$\hat{\tau} \in V_7$	Unit direction selected at modal sublimation; fixes the three quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$ that index the three matter generations	§1.7

A.2. Modal threshold and order parameter

Symbol	Meaning	Where
τ	Modal threshold value of the *global* tension $T(C)$; unsubscripted τ always means this one	§4.6
$\varepsilon = \lvert T(C)\rvert - \tau$	Excess tension (the control parameter)	Ch. 3
ψ	Modal order parameter (V_7-valued; the complex-scalar form is its 2-plane projection, cf. Ch. 3)	Ch. 3
$\psi^* = \pm\sqrt{2\varepsilon}$	Order parameter at the potential minimum	Ch. 4
$F[\psi, \varepsilon]$	Modal potential (Landau–Ginzburg form)	Ch. 4
$\mathcal{V}$	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	Ch. 4
r_0	Vacuum radius, $r_0^2 = \varepsilon/(2b)$	Ch. 4
m_0	PST elementary mass scale, $m_0 = 4\pi^3 r_0^2 d_0/(3c^2)$; the form is derived, the magnitude inherits d_0 and r_0	§9.5

A.3. Modal sublimation and projection

Symbol	Meaning	Where
$\mathbf{S}$	Precausal (Boolean / distinction) category; source of Φ	Ch. 4
$\mathbf{G}$	Geometric (manifold / spectral-triple) category; target of Φ	Ch. 4
$\Phi : \mathbf{S} \to \mathbf{G}$	Modal-sublimation functor (precausal → geometric)	Ch. 4
ρ	Realisation map from substrate to spacetime	Ch. 4
Π	Projection operator, substrate → four-dimensional field theory	Ch. 4

A.4. Foundational object: the real spectral triple

Symbol	Meaning	Where
$M = \mathbb{R} \times S^3$	Four-dimensional Lorentzian spacetime factor (KO-dimension 4)	§1.5
F	Finite noncommutative internal factor (KO-dimension 6)	§1.5
$M \times F$	Foundational real spectral triple (total KO-dimension $\equiv 2 \pmod 8$)	§1.5
(A, H, D, J, γ)	Spectral-triple data: algebra, Hilbert space, Dirac operator, real structure, chirality grading	§1.5
$A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$	Internal algebra of F	§1.5
$\mathrm{Im}(\mathbb{O}) \cong \mathbb{R}^7$	Imaginary octonions: the algebraic space on which F is built	Ch. 13
$\mathrm{Cl}(0,6)$	Clifford algebra on $\mathrm{Im}(\mathbb{O})$	Ch. 13
$\overleftarrow{\mathbb{O}} \cong \mathrm{Cl}(0,6) \otimes \mathbb{C} \cong M_8(\mathbb{C})$	Octonionic left-action algebra on $\mathbb{O}$ (Furey)	Ch. 13
$\omega = L_{e_1} \cdots L_{e_6}$	Cl(0,6) chirality grading volume element (Spin(6)-invariant)	Ch. 13
$\gamma_K = i\omega$	Substrate parity grading; equals the Cl(0,6) chirality grading in the Jordan–Wigner Majorana representation (Computation 19)	Ch. 13
$f = (I + i\omega)/2$	Chirality projector on H_F; its canonical rank-1 refinement is Furey's primitive idempotent	Ch. 13
KO-dimension	Connes mod-8 invariant of a real spectral triple; $4 + 6 = 10 \equiv 2 \pmod 8$ for $M \times F$. The name is from real K-theory, K for Grothendieck's *Klasse* and O for the orthogonal group a real bundle carries; it abbreviates no English phrase	§1.5

A.5. Standard Model and gauge sector

Symbol	Meaning	Where
$\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$	Standard Model gauge group; unimodular unitaries of A_F	Ch. 13
$\mathrm{SU}(3)_c$	Colour sector; the $M_3(\mathbb{C})$ factor of A_F	Ch. 13
$\mathrm{SU}(2)_L$	Left-chiral weak sector; the $\mathbb{H}$ factor with the directed modal threshold selecting the left factor of $\mathrm{SO}(4) \cong \mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$	Ch. 13
$N_c = 3$	Colour multiplicity; the rank of $M_3(\mathbb{C})$	Ch. 13
$N_{\mathrm{gen}} = 3$	Generation count; the three quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$ (§14.2)	Ch. 14
D_F	Yukawa / mass operator on H_F	Ch. 13
Y	$\mathrm{U}(1)_Y$ hypercharge	Ch. 13
$Z^2 \approx e^{-1}$	Field-normalisation ratio: substrate-side factor $Z_H(\beta_{\mathrm{KO}}) \to e^{-1}$ derived from P1–P3 (Computation 100); its matching to the SM-side ratio is bridge premise (B), whose standing the (B) row below records	§12.3

A.6. Scales and constants

Symbol	Meaning	Where
m_h	Higgs mass (125.25 GeV, the 2022 input used for M_*; the 2026 Review gives 125.13)	Ch. 7
v	Higgs vacuum expectation value (≈ 246 GeV)	Ch. 7
$M_* = 4\pi m_h \sqrt{2/3}$	Electroweak modal scale (≈ 1.285 TeV)	Ch. 7
$\ell_* = \hbar c / M_*$	Modal length scale ($\approx 1.535 \times 10^{-19}$ m)	Ch. 12

Symbol	Meaning	Where
d_0	Substrate coherence scale (Landau–Ginzburg coherence length)	Ch. 15
ℓ_P	Planck length	Ch. 8
M_P	Planck mass	Ch. 7
G	Newton's gravitational constant (consistency relation in PST)	Ch. 7
$\alpha_G \equiv G/d_0^2$	Gravitational-coupling coefficient (in $\hbar = c = 1$)	Ch. 8
$\xi = 90/\pi^2$	Casimir-correction coefficient (parameter-free under the Gaussian projection-kernel choice, dimensionless; the d^{-6} exponent is kernel-independent; distinct from the coherence length $\xi_\psi = \xi_{\mathrm{BCS}} = d_0$)	Ch. 15
Λ	Cosmological constant (Ch. 16) and, by reuse of notation, the spectral-action cutoff $\mathrm{Tr}\, f(D/\Lambda)$ (Ch. 7). The two roles are kept apart by context.	Ch. 16, Ch. 7

A.7. Substrate Dirac and Walsh sector (L_round closure)

Symbol	Meaning	Where
D_{sub}	Substrate Dirac operator on $\bigotimes_a \mathbb{C}^2$; built by Jordan–Wigner from the χ_a^{Cliff}, satisfying $D_{\mathrm{sub}}^2 = D \cdot I$	§18.3
χ_S	Walsh mode indexed by $S \subseteq D$; eigenvector of all χ_a^{Cliff}	§18.3
χ_a^{Cliff}	Clifford-valued Walsh generator for bit $a \in D$ (Jordan–Wigner)	§18.3

Symbol	Meaning	Where
T_{tail}	Walsh-tail operator $\sum_{\lvert S\rvert>k_D}\chi_S$; complement of the bridge support	§18.3
T_{tail}^{+-}	Off-diagonal block $P_+T_{\text{tail}}P_-$ in the D_{sub} eigenbasis	§18.3
σ_1,σ_2	Closed-form operator norms of T_{tail}^{+-} on S_D-isotypic components: $\sigma_1=2^{D-1}-D$ (trivial), $\sigma_2=D-2$ (standard rep)	§18.3
ρ_F	Fermionic-signed S_D representation acting on Walsh modes by $\rho_F(g)\lvert S\rangle=\operatorname{sgn}(g\lvert_S)\,\lvert g(S)\rangle$; commutes with D_{sub}	§18.3
k_D	Walsh truncation cutoff; the maximum $\lvert S\rvert$ retained in the bridge support	§18.3
L_round	Round-trip Lip-norm criterion (substrate → Bergman → substrate)	§18.3
L_comm	Commutator-norm Lip criterion: $\lVert[D_\alpha, b_C(\chi_S)\otimes I_2]\rVert/\lVert b_C(\chi_S)\rVert\to 2\sqrt{\lvert S\rvert}$	§18.3
L_F	Fréchet smooth Lip-norm $L_F(\cdot)=\sup_j\lVert\mathrm{ad}_D^j(\cdot)\rVert/j!$	§18.3

A.8. Bergman–Toeplitz sector and the refined bridge (L_comm closure)

Symbol	Meaning	Where
B^2	Open unit ball $\{(z_0,z_1)\in\mathbb{C}^2:\lvert z_0\rvert^2+\lvert z_1\rvert^2<1\}$	§18.3
$H^2_\alpha(B^2)$	Weighted Bergman space on B^2 with weight $(\alpha+3)(1-\lvert z\rvert^2)^\alpha/\pi^2$; canonical weight $\alpha=0$	§18.3
$H^2_\alpha(B^2)^{(N)}$	Truncation to monomials $z_0^a z_1^b$ with $a+b\le N$	§18.3

Symbol	Meaning	Where
$e_{a,b} = z_0^a z_1^b / \|z_0^a z_1^b\|$	Orthonormal monomial basis of $H^2_\alpha(B^2)$	§18.3
$w = z_0 z_1$	Diagonal monomial; the bridge symbols are polynomials in w	§18.3
T_f	Holomorphic Toeplitz operator with symbol f: $T_f h = P(fh)$, P = Bergman projection	§18.3
T_{f_k}	Refined 2-monomial bridge $T_{w^{m(k)}} + \alpha(k)\, T_{w^{m(k)+1}}$ at weight k	§18.3
J_+, J_-, J_z	SU(2) generators acting on holomorphic polynomials by z_0, z_1 rotation	§18.3
$D_\alpha = J_a \otimes \sigma_a$	SU(2)-equivariant round-S^3 Dirac on $H^2_\alpha(B^2) \otimes \mathbb{C}^2$	§18.3
$b_C(\chi_S) = T_{f_{\|S\|}}$	Bridge map: substrate Walsh mode $\to$ Bergman Toeplitz with the weight-class symbol	§18.3
$B(\cdot)$	Berezin transform: $B(T)(z) = \langle k_z, T k_z\rangle / \|k_z\|^2$, k_z = reproducing kernel	§18.3
N	Bergman truncation parameter; total-degree cutoff $a + b \le N$	§18.3
$\sigma_x, \sigma_y, \sigma_z$	Pauli matrices on the $\mathbb{C}^2$ spinor factor	§18.3
I_2	2×2 identity on the spinor factor	§18.3

A.9. Bridge parameters and closure-rate quantities (Lemmas 5–8)

Symbol	Meaning	Where
$m(k)$	Bridge monomial power for weight k; the unique integer with $m/2 < \sqrt{k}$ in the appropriate $(m, m+1)$ range (Lemma 5(b))	§18.3

Symbol	Meaning	Where
$\alpha(k)$	Closing coefficient from Lemma 5(c); pins the bulk ratio to $2\sqrt{k}$ exactly	§18.3
$\beta(k)$	Intermediate coefficient in the 3-monomial form of f_k (where used; absent in the 2-monomial form)	§18.3
q	Parametric variable in the Lemma 5(c) closed form $\alpha(q) = 2m(1-mq)/[(m+1)\sqrt{1-q^2}((m+1)q-1)]$	§18.3
$\mathrm{gap}(k, N)$	$\mathrm{ratio}(k, N) - 2\sqrt{k}$; finite-$N$ deviation from the bulk ratio	§18.3
$c_1(k), c_2(k)$	Leading and subleading coefficients of the gap expansion $\mathrm{gap}(k, N) = c_1/N + c_2/N^2 + O(1/N^3)$; $c_1(1) = -1$ exactly (Lemma 7(a))	§18.3
γ_D	L_comm closure rate at substrate size D: $\gamma_D = \max_{k \le k_D} \lvert\mathrm{gap}(k, N(D))\rvert$	§18.3

A.10. Substrate symmetry, complement, and configuration spaces

Symbol	Meaning	Where
$v : D \to V_7$	P1 direction-content assignment: each distinction a carries inherent V_7-valued direction $v(a)$; the third component of the substrate triple (D, δ, v), with $v(C) = \sum_{a \in C} v(a)$	Ch. 2
$\langle\cdot,\cdot\rangle$, $\times$: $V_7 \times V_7 \to V_7$	Octonionic inner and vector products on $V_7 \cong \mathrm{Im}(\mathbb{O})$; the bilinear product satisfies $\langle u \times v, u\rangle = 0$ and $\lvert u \times v\rvert^2 = \lvert u\rvert^2\lvert v\rvert^2 - \langle u, v\rangle^2$	Ch. 2

Symbol	Meaning	Where
$\bar{C} = D \setminus C$	Configuration complement; the set of distinctions absent from C	Ch. 2
$\Delta T(C) := T(C) - T(\bar{C})$	Net configuration-complement tension imbalance, a nonzero V_7-vector central to complement-asymmetry and the Bernoulli $p = 1/2$ derivation (distinct from the electroweak custodial ΔT, below)	Ch. 3
$\mathrm{Aut}(D, \delta)$	Distinction-preserving bijections of D; the canonical measure is invariant under them	Ch. 2
S_D, $S_D \times \mathbb{Z}_2$	Full symmetric group on D (driving ρ-equivariance); $S_D \times \mathbb{Z}_2$ adds complementation	Ch. 2, Ch. 6
S_3	Generation-permutation subgroup; forces the leading Yukawa degeneracy across the three generations	Ch. 13
$L^2(\mathcal{P}(D), \mu)$	PST quantum Hilbert space of square-integrable complex functionals on configuration space under μ	Ch. 4
$W^{1,2}(\mathcal{P}(D), \mu)$	Boolean Sobolev space $\{\psi : \mathbb{E}_\mu \lvert\nabla_{\mathcal{C}}\psi\rvert^2 < \infty\}$ on which the Dirichlet form lives	Ch. 4

A.11. Order parameter and the Boolean gradient

Symbol	Meaning	Where
$\hat{\eta} \perp \hat{\tau}$	Unit direction perpendicular to $\hat{\tau}$ spanning the projection 2-plane in V_7	Ch. 3
$\mathbb{C}_{\hat{\tau},\hat{\eta}}$	The 2-plane $\subset V_7$ spanned by $\hat{\tau}$ and $\hat{\eta}$ onto which the V_7 Landau–Ginzburg potential projects	Ch. 9

Symbol	Meaning	Where
$\psi_{\mathbb{C}}$	Complex-scalar order parameter $\psi_{\mathbb{C}} = \langle \psi, \hat{\tau} \rangle + i \langle \psi, \hat{\eta} \rangle$ recovering the textbook complex LG/Higgs slice	Ch. 9
$\hat{\psi} \in S^6$	Unit direction of the order parameter, $\hat{\psi} = \psi/\lvert\psi\rvert$ on the vacuum 6-sphere; modal sublimation locks it to $\hat{\tau}$	Ch. 3
$\nabla_{\mathcal{C}}\psi$, $\lvert\nabla_{\mathcal{C}}\psi\rvert^2$	V_7-valued Boolean Dirichlet gradient summed over single-bit-flip (Hamming) edges of $\mathcal{P}(D)$	Ch. 3
∂_a^B	Per-site Boolean difference, $\partial_a^B \psi = \psi(C \cup \{a\}) - \psi(C \setminus \{a\})$	Ch. 7
$C \triangle C'$	Symmetric difference of configurations; $\lvert C \triangle C' \rvert = 1$ indexes the gradient edges	Ch. 4

A.12. Mosco convergence and A7 reduction

Symbol	Meaning	Where
A1–A7	The seven explicit conditions for Mosco convergence; A7 (equicoercivity) is a theorem given A1, A2, A4 and A3, its equidistribution condition A3 supplied by P3's maximum-entropy sublimation, while A1 is architectural input	§18.1
(240)	The uniform-rate hypothesis to which A7 reduces; the K-uniform Berry–Esseen rate $\eta_{K,R} < d_0^2/4$	§18.1.6
$\eta_{K,R}$	Rate constant in the uniform-rate hypothesis; the equidistribution defect the continuum geometry can absorb while still closing	§18.1.6

C_K	Equicoercivity (closure-margin) constant, $C_K = (d_0^2/4 - \eta_{K,R})/(C_P(K)^2 + 1)$; $C_K > 0$ is the condition that spacetime closes as a manifold	§18.1.6
$C_P(K)$	Poincaré constant of compact $K \subset (M, g)$ for test functions supported in K	§18.1.6
D_n, ρ_n	Sequence of finite site sets $\lvert D_n\rvert \to \infty$ with realisation maps $\rho_n : D_n \to M$	§18.1
$\mathcal{T}_n$, ϕ_a	Delaunay triangulation of the realised cloud and its hat (Lagrange) basis functions	§18.2
a_n, V_M	Matched site spacing per scaling A1, $a_n \to 0$ with $\lvert D_n\rvert\, a_n^4 \to V_M$ the Riemannian volume of M (distinct from the fixed coherence length d_0)	§18.1
E^B / E_n^B	Boolean Dirichlet form on $L^2(\mathcal{P}(D_n), \mu_n)$; central object of the Mosco argument	§18.1
L^B / L_n^B	Boolean Laplacian; generator of E^B	§18.1
E^{Riem}, Δ_g	Mosco limits: the Riemannian Dirichlet form and the Laplace–Beltrami operator on (M, g)	§18.1
$\Pi^\dagger$ / $\Pi_n^\dagger$	Adjoint pullback $C_c^\infty(M) \to L^2(\mathcal{P}(D_n), \mu_n)$; comparison map in the Mosco recovery sequence	§18.1
$A_n(u)$	Empirical aggregate (exact pullback of E_n^B) whose mean deviation is the binding residual	§18.1.6
ξ_a	i.i.d. Gaussian per-site displacement with covariance $\sim d_0^2$	§18.2
$\mathcal{F}(K, R)$	Test-function class $\{u \in C_c^\infty(K) : \lVert u\rVert_{H^1} \le R\}$ on which the rate hypothesis is required uniform	§18.2
$c_{\mathrm{BE}} = 0.4748$	Optimal classical Berry–Esseen constant (Shevtsova), quoted only for the i.i.d. surrogate central-limit comparison; the proof's fluctuation channel is the McDiarmid concentration of Step 3	§18.2

ρ_3	Third absolute moment of the per-site displacement in the Berry–Esseen rate	§18.2

A.13. Higgs naturalness and the M_* one-loop derivation

Symbol	Meaning	Where
δm_h^2	One-loop quadratically-sensitive correction to the Higgs mass squared at cutoff M_*	§12.2
$\mathcal{V}(M_*)$	Veltman quadratic-sensitivity bracket $[6\lambda + \frac{9}{4}g^2 + \frac{3}{4}g'^2 - 6y_t^2] \approx -3.3$ at M_* (top-dominated, no zero crossing)	§12.2
$\Delta \equiv \lvert\delta m_h^2\rvert / m_h^2$	Fine-tuning (naturalness) parameter, ≈ 2.3, well within $\Delta \lesssim 10$	§12.2
$C_H = 6\lambda_{\mathrm{PST}} = \frac{3}{2}$	Higgs self-loop coefficient $(N+2)\lambda_{\mathrm{PST}}$, fixing $m_h^2 = C_H M_*^2/(16\pi^2)$	§12.2
$N_B = N_F = 2^{D-1}$	Substrate boson and fermion mode counts; their equality cancels the leading δm_h^2 divergence	§12.2

A.14. Partition-function / Chamseddine–Connes correspondence

Symbol	Meaning	Where
$\beta_{\mathrm{KO}} = 2$	Binary-gap tempering exponent $2 = 1 - (-1)$ (the spectral range of one binary distinction; the KO subscript is historical, not load-bearing)	§12.3
$Z_H(\beta)$	Tempered substrate partition function $Z(H_{\mathrm{Higgs}}, \beta) = \mathbb{E}_\mu[\exp(-\beta H_{\mathrm{Higgs}})] \to e^{-1}$ at β_{KO}	§12.5

Symbol	Meaning	Where
$H_{\text{Higgs}}(C) = \bar{X}$	Substrate Higgs Hamiltonian $= \lvert C\rvert/D$, forced by additivity, matched scaling, and uniformity (no free parameters)	§12.3
$\bar{X}(C) = \lvert C\rvert/D$	Empirical mean of bit-occupation over the substrate	§12.3
$B_a(C)$	Bit-occupation indicator for site a in configuration C	§12.3
F1–F11	The eleven labelled statements of the partition-function-level CC correspondence; rostered in full, with pages, in the next block	§12.5
Bridge Premise (B)	$\lambda_{\text{SM}}(M_*) = \lambda_{\text{PST}}^{\text{bare}} \cdot Z(H_{\text{Higgs}}, \beta_{\text{KO}})$; substrate side closed, the SM-side matching reduced to a field-strength-renormalisation identity (open)	§12.5
ϕ / Φ	SM-side Higgs field $\phi = \Pi(\psi)$ and doublet Φ (distinct from the substrate ψ and the functor Φ)	Ch. 12
$\Phi^\dagger\Phi$	Higgs bilinear $= \langle\psi, \psi\rangle$ under the projection	§1.7
μ^2	SM mass-squared parameter matched to PST excess tension ε	Ch. 12
λ_{SM}	Running SM quartic, $\lambda_{\text{SM}}(M_*) = b \cdot Z^2 \approx 0.092$; distinct from the substrate quartic $b = \lambda_{\text{PST}} = 1/4$	Ch. 12

A.15. The framework statements F1–F11

The eleven statements of the partition-function Chamseddine–Connes framework (§12.5), which the text cites by label alone.

Label	Kind and statement	§
F1	*Definition.* The PST partition-function spectral triple	12.6.1.1
F2	*Definition.* The β-partition function $Z(O,\beta) = \mathbb{E}_\mu[\exp(-\beta O)]$ of a substrate observable	12.6.1.2
F3	*Definition.* The matched-scaling exponent $\beta = 2$, the binary-distinction gap $2 = 1 - (-1)$	12.6.1.3
F4	*Proposition.* The PST effective theory below M_* is exactly the Standard Model. **The emergence premise**, and the one the closure ultimately leans on: the $\lambda_{\mathrm{SM}}(M_*)$ prediction is parameter-free only if it holds	12.6.2.1
F5	*Proposition.* Wilsonian matching at M_*	12.6.2.2
F6	*Proposition.* The PST UV Higgs quartic $\lambda_{\mathrm{PST}}^{\mathrm{bare}} = b = 1/4$	12.6.2.3
F7	*Proposition.* Substrate UV–IR scale separation at $\Lambda = \sqrt{D}$	12.6.2.4
F8	*Theorem.* The substrate Higgs Hamiltonian $H_{\mathrm{Higgs}} = \bar{X}$, forced by additivity, matched scaling and uniformity (Computation 89)	12.6.3.1
F9	*Theorem.* The binary-gap-tempered Bernoulli partition function $Z(H_{\mathrm{Higgs}}, \beta_{\mathrm{KO}}) \to \mathrm{e}^{-1}$, by the Cramér–Bernoulli asymptotic (Computation 88)	12.6.3.2
F10	*Theorem.* Unique soft mode; total reduction of the bridge premise	12.6.3.3
F11	*Theorem.* $Z^2 \to \mathrm{e}^{-1}$ in the substrate	12.6.3.5

A.16. Standard Model sector, extended, and the algebra synthesis

Symbol	Meaning	Where
$G_2 = \mathrm{Aut}(\mathbb{O})$	Exceptional 14-dimensional automorphism group of the octonions; preserves the inner and vector products on V_7	Ch. 13
$\mathrm{Stab}_{G_2}(\hat{\tau}) = \mathrm{SU}(3)_c$	Stabiliser of the threshold direction $\hat{\tau}$ in G_2; the source of colour $\mathrm{SU}(3)_c$ (Freudenthal–Tits)	Ch. 13
A_μ, $F_{\mu\nu}$	U(1) connection, massive here since the U(1) acting on θ is broken (the Z, Computation 172); the photon is the unbroken combination, with $D_\mu\theta = \partial_\mu\theta - A_\mu$, and its field strength $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$	§13.3
S_{Maxwell}	Gauge kinetic term, supplied by the spectral action's inner fluctuations (Computation 6); its coupling is not fixed by the theory	§13.3
$\theta(x)$	Local Goldstone / vacuum phase assigned per spacetime point by Π, gauging the global U(1)	§13.3
$H = (H^+, H^0)^T$	SM Higgs doublet; the neutral component H^0 is identified with ψ in the broken phase	§13.4
$\mathcal{V}_{\mathrm{EW}}$	Electroweak vacuum manifold $\{H : \lvert H\rvert^2 = v^2/2\} \cong S^3 \cong \mathrm{SU}(2)$	§13.4
$\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$	Isometry group of $\mathcal{V}_{\mathrm{EW}}$; $\mathrm{SU}(2)_R$ broken by the directed threshold crossing, leaving chiral $\mathrm{SU}(2)_L$	§13.4
λ_a, f_{abc}	Gell-Mann matrices and their structure constants in the octonionic colour derivation	Ch. 13
C_2	SU(3) quadratic Casimir $\sum_a (\mathrm{ad}\,\Lambda_a)^2/4$; irrep discriminator in the $M_8(\mathbb{C})$ decomposition	Ch. 13

α_k	Fermionic ladder operators $\frac{1}{2}(L_{2k-1} + iL_{2k})$ from octonionic left-multiplications, building one generation's colour triplet	Ch. 13
Q_L, u_R, d_R, L_L, e_R, ν_R	One-generation SM fermion multiplets: quark doublet, up/down singlets, lepton doublet, charged-lepton singlet, right-handed neutrino	Ch. 13
$\tilde{H} = i\sigma_2 H^*$	Conjugate Higgs ($Y = -1$) for the up-type Yukawa coupling	Ch. 13
$c_a, c_a^\dagger$, n_a	CAR creation/annihilation operators with $\{c_a, c_b^\dagger\} = \delta_{ab}$ and occupation $n_a = c_a^\dagger c_a \in \{0, 1\}$	Ch. 11
$N = \sum_a n_a$, $\mathcal{H}_D = \bigoplus_k \wedge^k \mathbb{C}^{\lvert D\rvert}$, $\mathcal{H}_D^\pm$	Total occupation grading the exterior-algebra Fock space $\mathcal{H}_D$ by $(-1)^N$ into bosonic/fermionic sectors $\mathcal{H}_D^\pm$	Ch. 11
$T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$	Dixon hyperspinor algebra (64-dimensional); the full electroweak setting (distinct from the tension functional T)	Ch. 13
$\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$	Local Lorentz Clifford algebra of M; its right-$\mathbb{H}$ action supplies the internal SU(2)	Ch. 13
$\mathcal{S}$	Minimal left ideal $(\mathrm{Cl}(0,6) \otimes \mathbb{C})\, p \cong \mathbb{C}^8$, generated by f's canonical rank-1 refinement p; one generation's content	Ch. 13
$\mathcal{C}$	Furey chain algebra; its *commutant* on the minimal left ideal is $\mathbb{C}$ (the overall phase), and the algebra *generated by* $\mathrm{Stab}_{G_2}(e_7)$ on the chiral half is $\mathbb{C} \oplus M_3(\mathbb{C})$	Ch. 13

A.17. Spectral action, Seeley–DeWitt, and cutoffs

Symbol	Meaning	Where
f_0, f_2, f_4	Cutoff-function moments $f_0 = f(0)$, $f_2 = \int_0^\infty u f(u)\,du$, $f_4 = \int_0^\infty u^3 f(u)\,du$ of the spectral-action heat-kernel expansion	§7.8
κ_{grav}	$O(1)$ scheme constant in $1/(16\pi G) = (f_2\Lambda^2/\pi^2)\,\kappa_{\mathrm{grav}}\,N_{\mathrm{dof}}$	§7.8
$a_1 = \frac{1}{6}$	Universal first Seeley–DeWitt coefficient, forcing the leading R (Einstein–Hilbert) curvature term	§7.8
$\mathcal{A}_2(x)$	Second Seeley–DeWitt coefficient ($\frac{1}{180}(\dots) + \frac{1}{72}R^2$ curvature combination) in the gravitational spectral-action expansion	§7.8
$\Lambda_{\mathrm{UV}} = 1/\ell_*$	Physical UV cutoff; a genuine breakdown scale of the continuum description, not a regulator	Ch. 13
$\rho_{\mathrm{A}}^{\mathrm{bare}}$	Bare Goldstone zero-point density $\Lambda_{\mathrm{UV}}^4/(16\pi^2)$ + counter-terms at fixed UV regulator	§9.1
$\Lambda_{\mathrm{eff}} = M_*/\sqrt{c}$	Effective four-fermion / dimension-six operator scale (13 to 16 TeV under the universal-kernel reading Computation 205 excludes; far above the LHC bound in the surviving reading), with c the (loop-suppressed) $O(1)$ dimension-six Wilson coefficient	§15.11
$\Lambda = \sqrt{D}$	Matched-scaling substrate spectral-action cutoff, at which the substrate Dirac eigenvalues collapse to ± 1	Ch. 12

A.18. Quantum mechanics from the substrate

Symbol	Meaning	Where
$\hat{x}_\mu$	Position operator, $(\hat{x}_\mu\psi)(x) = x_\mu\,\psi(x)$ on $L^2(M, d^4x)$	§10.2
$\hat{p}_\mu = -i\hbar\,\partial_\mu$	Momentum operator; the strong limit of the projected Boolean gradient $\Pi(\partial^B_\mu)$	§10.3
P_i	Born-rule orthogonal projectors onto outcome eigenspaces, $P_iP_j = \delta_{ij}P_i$, $\sum_i P_i = \mathbf{1}$	§10.2
$\mathcal{S}_i$	Measurement partition cell $\subseteq \mathcal{P}(D)$ whose projected wavefunction lies in eigenspace $\mathcal{H}_i$; the $\{\mathcal{S}_i\}$ partition $\mathcal{P}(D)$	§10.3
$\tau_i(C)$	Outcome-dependent modal threshold $\tau_0 + \kappa(\lambda o_i)^2$ crossed at measurement (distinct from the modal threshold τ)	§10.3

A.19. Cosmology and equation of state

Symbol	Meaning	Where
ρ_{inst}	Effective energy density of ongoing instantiation, $\rho_{\text{inst}} = j_\tau\,\Delta\mathcal{F}/c^2 > 0$; $\dot{\rho}_{\text{inst}} = 0$ (does not dilute)	§16.5
$w = -1$	Equation-of-state parameter: the instantiation reservoir's pressure over its energy density, $P_{\text{inst}}/\rho_{\text{inst}}$, with $P_{\text{inst}} = -\rho_{\text{inst}}$ forced by covariant energy conservation at $\dot{\rho}_{\text{inst}} = 0$; a structural theorem (de Sitter, not anti-de Sitter)	§16.5
$\Delta\mathcal{F}$	Free-energy (tension) released per threshold crossing, $\mathcal{F}_{\text{pre}} - \mathcal{F}\|_{\mathcal{V}} > 0$	§16.5
j_τ	Instantiation current; threshold-crossing rate, spatially uniform by the Cosmological Principle theorem and constant in time by non-dilution	§16.5

Symbol	Meaning	Where
$\Lambda_{\mathrm{PST}} = 8\pi G\rho_{\mathrm{inst}}/c^2$	Positive effective cosmological term contributed by ongoing sublimation	§16.5
Λ_{obs}	Observed cosmological constant in curvature units ($\approx 1.1 \times 10^{-52}\,\mathrm{m}^{-2}$; the Λ of equation (48) is the vacuum energy density, $8\pi G$ times smaller in these units); its magnitude remains $T(C)$-contingent and open	§16.5

A.20. Overloaded symbols (disambiguation)

Symbol	Meaning	Where
κ	Overloaded: scalar curvature (in Φ); substrate-apparatus measurement coupling (collapse rule); and the spectral-action moment factor κ_{grav}. Kept apart by context	Ch. 4, §10.3, §7.8
ΔT	Overloaded: the substrate complement-tension $\Delta T(C) = T(C) - T(\bar{C})$ (a V_7-vector) versus the electroweak custodial / oblique T-parameter ΔT ($\Delta T = 0$ structurally). Distinct objects	Ch. 3, Ch. 7
τ family	Three distinct thresholds, kept apart by decoration: bare τ is the modal threshold on the global tension $T(C)$; τ_{null} is the null threshold on the *pairwise* tension $T(\{a, b\})$, the causal boundary supplying the Lorentzian signature; τ_i is the modal threshold shifted by the measurement coupling. $\hat{\tau}$ is a direction, not a value	§4.6, §4.10, §10.3

Symbol	Meaning	Where
d_0 vs ℓ_*	**Two different lengths, ten orders apart at the bound and one under the canonical normalisation, and the one confusion this appendix most needs to prevent.** $d_0 \lesssim 2.3$ nm is the *coherence length* of the modal vacuum, an LG-dynamics quantity that sets the Casimir amplitude. $\ell_* = \hbar c/M_* \approx 1.5 \times 10^{-19}$ m is the substrate *grain*, a momentum cutoff. The ultraviolet cutoff is $\Lambda_{\mathrm{UV}} = 1/\ell_*$, never $1/d_0$; reading it off d_0 would put it at tens of electronvolts instead of $M_* \approx 1285$ GeV, and Computation 167 prices the difference at $(d_0/\ell_*)^6 \sim 10^{61}$ at the bound; under the canonical normalisation $d_0 = 2.2$ am and the ratio is 14.5 (Computation 201).	§15.3.1
$\mathcal{V}$	Overloaded by an explicit convention: $\mathcal{V} \cong S^6 \subset V_7$ is the full vacuum manifold, while $\mathcal{V} \cong S^1$ means its projection to the 2-plane slice, written $\mathcal{V}_{\hat{\eta}} \cong S^1$ where the distinction matters, as declared at §5. Appendix C uses both senses.	§5
τ_i	Overloaded: the outcome-dependent modal threshold $\tau_i(C)$ of the collapse rule, versus the self-adjoint exchange involution τ_i of the register rows for Computations 88, 100 and 112	§10.3, Appendix B
γ_D	Overloaded: the L_comm closure rate of §18.3 versus the L_round rate $\gamma_D \approx 0.37\,\mathrm{e}^{-0.28D}$ of the register rows for Computations 20, 23 and 51; the text names which	§18.3, Appendix B

Symbol	Meaning	Where
$\mathcal{S}$, $\mathbf{S}$, $\mathcal{S}_i$	Overloaded: $\mathbf{S}$ is the category of above-threshold configurations in the functor $\Phi : \mathbf{S} \to \mathbf{G}$; $\mathcal{S}$ is also the minimal left ideal of the chain algebra in the fermion chapters, and $\mathcal{S}_i$ a measurement cell of the collapse rule	§5, §14.2, §10.3
N	Overloaded: the Bergman truncation of the bridge in Chapter 18, the total occupation number of the Fock space, and an $O(N)$ count in the register	§18.3, Appendix B
w	Overloaded: the equation-of-state parameter $w = p/\rho$ of the cosmology chapter, the product $z_0 z_1$ in the two-site examples of the register, and a weight vector $w = (w_i)$ in the Born-rule example of §10.3.2.5	§16.5, Appendix B
D	Overloaded: the substrate's property set D, with $\lvert D\rvert$ its size and 2^D its configuration count, and the Dirac operator in the spectral triple (A, H, D, J, γ); the triple's D carries its name in the text	§1.5, §18.1

Appendix B: Computation Scripts

Each computation cited in the body is checked symbolically or numerically by the Python script listed below. Every script can be run in the browser, or downloaded, at the companion site, meelker.nl/pst-se/computations, where it appears as `computation_NN_se.py`; the longest, Computations 64 and 150, take several minutes, and every quoted number is the script's own output. The register is the complete research record, not only the cited subset: a handful of entries are exploratory probes or routes that a later computation superseded, and the text does not cite them. They are kept because the record of what was tried, and abandoned, is part of the account.

Each result carries one *provenance* tag recording the character of its verification: [**PROOF**] a rigorous proof or exact symbolic/algebraic construction; [**NUMERICAL**] a numerical or finite-size check from which an asymptotic is inferred; [**SURROGATE**] a result established for an i.i.d. surrogate or toy model standing in for a not-yet-proven emergence-determined quantity; [**EMPIRICAL**] an empirical or observational bound, renormalisation-group equation (RGE) test, or comparison to data. A result is current unless its tag notes that a later computation supersedes (or partially retracts) it; readers weighing rigour should consult the provenance tag and the description. Every [**SURROGATE**] row ends by naming what its surrogate stands in for, and the consistency audit requires it; a route tried on a framing later abandoned is tagged by what it computed, not as a surrogate.

Computation 1 · §14.3 · [PROOF]
Yukawa hierarchy.

Generation-symmetry structural-scope 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 $D = 6$).

Computation 2 · §4.12 · [PROOF]

Lorentzian signature.

Lorentzian signature $(+, -, -, -)$ from Derrick-style hyperbolic well-posedness of the Goldstone wave equation; the macroscopic dimension $n = 4$ is fixed by the minimal Connes-spectral-triple split given the verified substrate KO-6 (Computation 3) and the total KO-2 of the product.

Computation 3 · §1.5 · [PROOF]

Explicit product spectral triple of KO-dimension 2.

Builds the explicit product spectral triple (A, H, D, J, γ) for $M \times F$ where M is the 4D Lorentzian spacetime (KO-4) and F is the substrate internal factor (KO-6), verifying total KO-dim $4 + 6 = 10 \equiv 2 \pmod 8$, the Connes Standard-Model value whose sign $J^2 = -1$ makes the Euclidean fermionic action well-defined. Five checks: (K1) explicit KO-4 representative triple for M with signs $(-1, +1, +1)$; (K2) two substrate carriers of the KO-6 internal factor F, the octonion $\mathbb{C}^8$ picture ($\gamma_F = i\, L_{e_1} \cdots L_{e_6}$, $J_F = K$) and the bit-space complementation/parity triple at $D = 3$, both deliver KO-6 $(+1, +1, -1)$, so the two F-pictures agree at the KO level; (K3) F is even and carries colour: $[\gamma_F, \mathfrak{su}(3)] = 0$; (K4) the explicit 32-dim (γ, J, D) for $M \times F$ with a generic admissible internal D_F is verified to be a genuine real spectral triple of KO-2, so $4 + 6 = 10 \equiv 2$ is a real product of spectral triples, not just mod-8 bookkeeping; (K5) the two open inputs (D_M from spin structure, D_F as Yukawa/order-one selection) are explicitly located. Everything else (A, H, γ, J, the product law) is in hand.

Computation 4 · §7.8 · [PROOF]
Spin structure on $\mathbb{R} \times S^3$.

Closes the spin-structure step for $M = \mathbb{R} \times S^3$, splitting into a topological part (closed rigorously) and an analytic part (scoped). *Topological closure*: every orientable 3-manifold is parallelisable (Stiefel), in particular $S^3 = \mathrm{SU}(2)$; the spin obstruction $w_2(S^3) \in H^2(S^3; \mathbb{Z}/2) = 0$ vanishes; spin structures form a torsor over $H^1(S^3; \mathbb{Z}/2) = 0$, so the spin structure is unique; globally hyperbolic $M = \mathbb{R} \times \Sigma$ is spin iff Σ is spin, so $\mathbb{R} \times S^3$ carries a canonical Dirac operator. *Analytic part (scoped, residual)*: Mosco / strong-resolvent convergence is naturally a statement about Dirichlet forms / Laplacians; on a spin manifold the Dirac operator is the canonical first-order square root of (a curvature-shifted) Δ (Lichnerowicz: $D^2 = \nabla^*\nabla + R/4$). With the spin structure fixed and unique, D_M is determined; the remaining task is to show the substrate's Boolean→CAR/Clifford structure converges to this canonical D_M, not merely that Δ converges. Numerical verification of the canonical S^3 Dirac spectrum and its Weyl law (spectral dimension 3, eigenvalues $\pm(n + 3/2)/\ell$) is included.

Computation 5 · §8.5 · [NUMERICAL]
Einstein–Hilbert and Newton's G from the spectral action.

Derives G structurally from the Chamseddine–Connes spectral action $S = \mathrm{Tr}\, f(D/\Lambda)$, replacing the Kaluza–Klein dimensional-reduction identity (moot in the 10-foundational reading). The 4D heat-kernel asymptotic expansion (Chamseddine–Connes 1996; Chamseddine–Connes–Marcolli 2007) reads $S \sim 2f_4\Lambda^4 a_0 + 2f_2\Lambda^2 a_2 + f_0 a_4 + O(\Lambda^{-2})$ with moments $f_2 = \int_0^\infty u f(u)\, du$, $f_4 = \int_0^\infty u^3 f(u)\, du$, $f_0 = f(0)$. The three terms: $\Lambda^4 a_0 \to$ cosmological constant + volume; $\Lambda^2 a_2 \to$ Einstein–Hilbert $(1/16\pi G)\int R\sqrt{g}\, d^4x$; $\Lambda^0 a_4 \to$ Yang–Mills + Higgs + Weyl^2 (the SM bosonic sector). So gravity is neither put in

by hand nor obtained via Kaluza–Klein reduction; the EH term is the Λ^2 coefficient. Verifies the cutoff moments f_0, f_2, f_4 are finite positive for a sample f ($f(u) = e^{-u}$); the closed-form scaling $1/(16\pi G) = (f_2\Lambda^2/\pi^2)\,\kappa_{\text{grav}}\,N_{\text{dof}}$ with κ_{grav} an $O(1)$ scheme constant and N_{dof} the fermionic multiplicity, solving for Λ given observed G and the SM fermion count places Λ at a unification/Planck-ish scale.

Computation 6 · Chapter 13 · [PROOF]

Gauge group from A_F via unimodular unitaries; gauge bosons and Higgs as inner fluctuations.

Derives the Standard-Model gauge group $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ as the unimodular unitary group of $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ via the Chamseddine–Connes machinery, with the gauge fields arising as inner fluctuations of the Dirac operator $D \to D + A + \varepsilon' JAJ^{-1}$, $A = \sum_i a_i[D, b_i]$ for $a_i, b_i \in A_F$. This replaces the G_2-holonomy route to SU(3) (which dies in the 10-foundational reading) with the standard Connes derivation applied to the substrate-supplied A_F. Four checks: (N1) Lie-algebra dimensions $u(A_F) = 13$; unimodular condition removes one dim $\to 12 = \dim(\text{SM group})$; (N2) $\mathfrak{su}(2)$ from imaginary quaternions, $\mathfrak{su}(3)$ from Gell-Mann generators, both closure checks exact; (N3) inner-fluctuation count: gauge-boson count = adjoint = 12; spacetime $[D_M, a]$ gives the vector bosons A_μ, the finite $[D_F, a]$ gives the Higgs; (N4) U(1) bookkeeping: which U(1) is removed by unimodularity, which survives as $\mathrm{U}(1)_Y$.

Computation 7 · §14.4 · [PROOF]

Order-one condition selects the Yukawa form of D_F on the chiral bimodule.

Constructs an explicit one-generation lepton-sector finite spectral triple (8-

dim: ν_L, e_L, ν_R, e_R + 4 antiparticles) and verifies the chirality and order-one structure. Track J established that the naive substrate ladder SU(2) is, under either natural chirality grading, vector-like or chirality-mixing, *not* the SM's chiral $\mathrm{SU}(2)_L$. The Connes resolution uses the CHIRAL representation: $\mathbb{H}$ acts on left-handed doublets only, so $\mathrm{SU}(2)_L$ is chiral by construction; the order-one condition $[[D_F, a], Jb^*J^{-1}] = 0$ then SELECTS the Yukawa form of D_F. Four checks: (O1) the representation is even ($[\gamma, a] = 0$) and the $\mathrm{SU}(2)_L$ action of $\mathbb{H}$ is chiral, supported on left-handed particles only; (O2) order-zero $[a, Jb^*J^{-1}] = 0$ for all $a, b \in A_F$; (O3) order-one $[[D_F, a], Jb^*J^{-1}] = 0$ for the Yukawa D_F, and FAILURE for a generic admissible D_F, so order-one is a genuine constraint that selects the Yukawa form; (O4) the KO-6 sign of the finite real structure ($\varepsilon = +1$), consistent with KO-2 product. Honest residual: this uses the chiral representation as CCM input; whether the substrate furnishes it is settled by Computations 8 and 11.

Computation 8 · §14.4 · [PROOF]

Chirality bridge: γ_M vs the directed modal threshold.

Tests the two candidate substrate mechanisms for furnishing the chiral bimodule of Computation 7, locating chirality at the directed modal threshold rather than at γ_M. (a) γ_M does *not* furnish weak chirality: the physical chirality of an internal generator $1 \otimes T$ is governed entirely by $[T, \gamma_F]$; the γ_M factor tensors along and cannot change it (negative finding, structurally located). (b) The substrate DOES offer $\mathrm{SU}(2)_L \times \mathrm{SU}(2)_R$: quaternion left- and right-multiplications give two commuting $\mathfrak{su}(2)$'s $\cong \mathfrak{so}(4)$, the S^3 isometry algebra. Chirality selection is the choice of one factor, and the directed modal threshold $\varepsilon = |T(C)| - \tau > 0$ (book Ch. 13 parity-violation mechanism) is the physical principle that selects between the parity-conjugate choices. (c) Residual: the selected $\mathrm{SU}(2)_L$ must act on the $\gamma_F = +1$ block only (bimodule support); bridged here to the threshold mechanism and closed in Computation 11.

Computation 9 · §18.3 · [NUMERICAL]

Dirac convergence scoping.

Dirac-convergence scoping in spectral propinquity; numerical companion verifies $\mathrm{Cl}(0, 2D)$ 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 (256) (the full first-order route; the committed closure is the relaxed criterion of §18.3, whose all-(k, N) uniform bound is the surviving quantitative refinement).

Computation 10 · §15.10 · [EMPIRICAL]

Empirical M_* bounds.

SMEFT operator-by-operator survey of $M_* = 4\pi m_h\sqrt{2/3}$ at PST's tree level vs LHC/EWPO bounds; $\Delta S_{\mathrm{PST}} = \Delta T_{\mathrm{PST}} = 0$ structurally. *Forward reference*: the script's $M_* = 1.573$ TeV predates the off-by-one correction, and its Wilsonian reading with $O(1)$ coefficients is the universal-kernel reading Computation 205 excludes; its structural argument is the Higgs-doublet reading that survives there.

Computation 11 · §14.4 · [PROOF]

Bimodule support: threshold-selected $\mathrm{SU}(2)_L$ acts on the $\gamma_F = +1$ block only.

Closes the chiral-bimodule residual left open by Computation 8. Constructs a minimal but representative chiral H_F: one weak doublet $((\nu_L, e_L), (\nu_R, e_R))$ with $\gamma_F = \mathrm{diag}(+1, +1, -1, -1)$. Two representations of A_F are mathematically possible *a priori*: Choice A ($\mathbb{H}$ acts on $H_F^{(+)}$ left block, trivial on $H_F^{(-)}$, the parity-violating observed SM); Choice B ($\mathbb{H}$ acts on $H_F^{(-)}$ right block, triv-

ial on $H_F^{(+)}$, the parity-conjugate model, never observed). Both give consistent spectral triples in the Connes–Chamseddine sense; the chirality of $\mathrm{SU}(2)_L$ therefore reduces to "which of A or B is realised in nature?". PST's claim is that the directed modal threshold (book Ch. 13) is the physical principle that selects Choice A. Verifies (i) both representations of $\mathbb{H}$ commute with γ_F (necessary condition); (ii) both have clean block support; (iii) the substrate-side $\mathfrak{su}(2)_L$ and $\mathfrak{su}(2)_R$ operator-norm signatures under the threshold $\varepsilon > 0$ break the L↔R symmetry as required, selecting Choice A.

Computation 12 · §13.6 · [PROOF]

Consistency of octonionic-SU(3) embedding with Computations 1–11.

Examines whether the octonionic-SU(3) embedding ($\mathrm{SU}(3) \subset G_2 = \mathrm{Aut}(\mathbb{O})$ acting as automorphisms preserving the octonion product) is consistent with the rest of PST's structure, vs. the canonical block embedding ($\mathrm{SU}(3) \subset \mathrm{SU}(4) \cong \mathrm{Spin}(6)$ acting as block $\mathrm{diag}(\mathrm{SU}(3), 1)$ on the half-spinor). If the octonionic interpretation is consistent with Computations 1–11, Furey's three-generation construction ports directly into PST without changing framework assumptions. Six structural checks: (§1) two SU(3) subgroups of Spin(6) compared; (§2) Computation 6's gauge group derivation does not specify which SU(3); (§3) Computation 7's order-one transfers under either embedding; (§4) inner fluctuations and the Higgs sector are unaffected; (§5) Computation 4's spin structure and Computation 9's convergence are independent of the embedding; (§6) verdict: octonionic interpretation is structurally consistent with PST's existing results.

Computation 13 · §13.6 · [PROOF]

Canonical vs octonionic SU(3) **side-by-side comparison.**

Computes the consequences of canonical vs octonionic SU(3) embedding by re-deriving Computation 1, section 7 of the script (inner-fluctuation extension of the Yukawa structural-scope theorem) and Computation 7 (order-one condition selecting Yukawa form) under both, then compares for "best PST fit". Best-fit criteria in priority order: (C1) does $N_{\text{gen}} = 3$ emerge structurally? (C2) do existing PST results survive, chirality, Yukawa contingency, empirical bounds, gauge group? (C3) is the construction algebraically clean, no new posits? (C4) consistency with the Connes–Chamseddine framework? Seven sections culminate in the recommendation of the octonionic embedding as the natural source of $N_{\text{gen}} = 3$ via Furey's six-SU(3)-triplet decomposition of $\text{Cl}(0, 6)$, with the open work being the explicit chain-algebra construction (Computations 14–17).

Computation 14 · §13.6 · [PROOF]

Inner-fluctuation structural-scope theorem extension (Computation 1, section 7 of the script) under octonionic embedding.

Step B of the octonionic-SU(3) re-derivation programme identified by Computation 13. Original Computation 1, section 7 of the script (canonical embedding) assumed $H_F = H_F^{(1)} \otimes \mathbb{C}^{N_{\text{gen}}}$ tensor-factor structure for generations, with A_F acting as $a^{(1)} \otimes I_{\text{gen}}$ (generation-blind), so $[D_F, b]$ inherits generation structure of D_F unchanged. Under the octonionic embedding, H_F is *not* in tensor-factor form for generations; the three "generation copies" are internal states of the 64-complex-dim algebra $\mathbb{C} \otimes \overleftarrow{\mathbb{O}}$, with A_F's $M_3(\mathbb{C})$ acting via the octonionic SU(3), mixing generation copies by SU(3) rotations. The canonical §7 argument "A_F is gen-blind" does *not* transfer. The substantive question is whether the structural-scope STATEMENT itself still holds:

can inner fluctuations turn a gen-symmetric $D_F^{(0)}$ into one with generation-asymmetric Yukawa eigenvalues? Five sections: (§1) toy octonionic-SU(3) model; (§2–§3) inner fluctuations under each embedding; (§4) substrate-level S_3 constraint as the deeper invariant; (§5) verdict: individual fluctuations can mix generations but substrate S_3 invariance constrains physical observables; Yukawa contingency preserved.

Computation 15 · §13.6 · [NUMERICAL]

Order-one Yukawa selection (Computation 7) under octonionic embedding.

Step C of the octonionic re-derivation programme. Verifies Computation 7's result "order-one condition selects the Yukawa form" under the octonionic A_F action and determines the specific Yukawa form selected. Original Computation 7 (canonical): H_F = 8-dim lepton-sector finite triple, A_F acts via canonical lepton-sector representation ($M_3(\mathbb{C})$ trivial on leptons, $\mathbb{H}$ on left doublet), order-one $[[D_F, a], Jb^*J^{-1}] = 0$ SELECTS the Yukawa form of D_F. Under the octonionic embedding, A_F's action on H_F changes (the chain-algebra structure of $\overleftarrow{\mathbb{O}}$), so the order-one condition is a different algebraic constraint. Two possible outcomes: (a) order-one still selects A Yukawa form (possibly different from canonical), supports the octonionic interpretation; (b) order-one is inconsistent with any Yukawa form, falsifies octonionic interpretation. Seven sections culminate in the verdict: order-one selecting condition is unitarily covariant; Yukawa residue $\sim 2 \times 10^{-5}$ vs ~ 10 for generic D_F under BOTH embeddings, the selecting condition transfers, and the octonionic interpretation passes this test.

Computation 16 · §13.5 · [NUMERICAL]
Furey foundation: $\overleftarrow{\mathbb{O}} \cong \mathbb{C} \otimes \mathrm{Cl}(0,6) \cong M_8(\mathbb{C})$, octonionic SU(3) as G_2-stabiliser.

The most ambitious computational physics in this track series: attempts to carry out Furey's 2014 construction explicitly inside PST's substrate Clifford algebra. Twenty-two sections; the core run: (§1) octonion multiplication table (Fano-plane convention); (§2) left-multiplication operators L_{e_a} as 8×8 complex matrices; (§3) verify the algebra structure of $\overleftarrow{\mathbb{O}}$ as 64 explicit 8×8 matrices; (§4) identify the primitive idempotent f (vacuum) of $\overleftarrow{\mathbb{O}}$; (§5) construct the octonionic $\mathrm{SU}(3) \subset G_2$ as stabiliser of f and verify the 8 adjoint elements with exact commutation relations; (§6) decompose $\overleftarrow{\mathbb{O}}$ (64-complex-dim) under this SU(3) into irreducible representations; (§7) compare to Furey's $16 + 48$ split; (§8) honest assessment of how rigorously this verifies the three-generation interpretation; (§9–§22) extended analysis, including the convention-aligned complex structure $J = L_{e_7}$ with Cayley pairing $(e_1, e_3), (e_2, e_6), (e_4, e_5)$ (§16) and the full SU(3) lift. Verifies $\overleftarrow{\mathbb{O}} \cong \mathbb{C}\otimes\mathrm{Cl}(0,6) = M_8(\mathbb{C})$ and octonionic SU(3) as G_2-stabiliser of e_7 to machine precision. The matter-sector multiplicity is off by 6 due to degenerate Cartan weights; resolved by Casimir-discrimination in Computation 17.

Computation 17 · §13.5 · [PROOF]
Definitive SU(3) decomposition of $M_8(\mathbb{C})$ via Casimir.

Closes the off-by-6 residue from Computation 16 by using the SU(3) quadratic Casimir $C_2 = \sum_a (\Lambda_a)^2$ as an irrep discriminator. Computation 16 identified the eight-generator octonionic SU(3) subgroup (8 adjoint elements verified, commutation relations exact), but the matter-sector multiplicity count was off because the joint (Λ_3, Λ_8) eigenbasis is not unique on degenerate weights, singlets and $(0,0)$ members of triplets share Cartan weights and can only

be told apart by the Casimir. Six sections: (§1) reconstruct the octonionic SU(3) generators (lifted from Gell-Mann via the $J = L_{e_7}$ complex structure on $(e_1, e_3), (e_2, e_6), (e_4, e_5)$, the correct Cayley pairing for the G_2-stabiliser of e_7); (§2) build $\mathrm{ad}(\Lambda_a)$ on $M_8(\mathbb{C})$; (§3) build the Casimir C_2 as a 64×64 Hermitian matrix; (§4) diagonalise C_2 and group eigenvalues by SU(3) irrep Casimir values (0 for **1**, 4/3 for **3**, 3 for **8**, 10/3 for **6**); (§5) confirm the decomposition $6 \cdot \mathbf{1} + 5 \cdot \mathbf{3} + 5 \cdot \overline{\mathbf{3}} + 2 \cdot \mathbf{8} + \mathbf{6} + \overline{\mathbf{6}}$ (total dim 64) matches the Clebsch–Gordan expectation to machine precision; (§6) status statement for the octonionic-SU(3) item: structurally derived, decomposition exact.

Computation 18 · §13.7 · [PROOF]

ω as the unique Spin(6) singlet.

Spin(6)-invariance in $\mathrm{Cl}(0,6) \otimes \mathbb{C}$: $\omega = L_{e_1} \ldots L_{e_6}$ 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) non-trivial Spin(6)-invariant idempotent: the chirality projector, whose canonical rank-1 refinement is Furey's primitive idempotent.

Computation 19 · §13.7 · [PROOF]

Substrate parity grading = $\mathrm{Cl}(0,6)$ chirality grading.

Schur on the irreducible chiral halves $H_\pm$ of $H_F = \mathbb{C}^8$ gives Hermitian Spin(6)-commutant = span$\{I, \gamma\}$ (2-real-dim). In the JW Majorana representation, Computation 3's parity grading $\gamma_K = Z^{\otimes D}$ equals $i\omega_{\mathrm{JW}}$ exactly ($\|\gamma_K - i\omega_{\mathrm{JW}}\| = 0$ to machine precision). The chirality projector $f = (I + i\omega)/2$ is delivered directly by the substrate's parity-selected real structure; Furey's primitive idempotent is its canonical rank-1 refinement [121], and her acting-on-$\mathbb{C} \otimes \mathbb{O}$ chain-algebra construction then delivers $\mathrm{SU}(3) \times \mathrm{U}(1)$ on one generation under $\mathrm{SU}(3)_c$. The synthesis to the full $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ requires the $\mathbb{H}$ factor from the emergent Lorentzian Dirac sector $\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ (Computation 66)

combined via Dixon's hyperspinor framework $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ [146].

Dirac-convergence step (§18.3): L_round and L_comm computations.

Computation 20 · §18.3 · [NUMERICAL]
L_round sweep.

PST analogue of [119] Lemma 3.2 across $D = 4, 6, \ldots, 20$ 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.

Computation 21 · §18.3 · [PROOF]
Dirac commutator Lip-norm scaling.

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 [113]. The first-order Dirac structure on the substrate is polynomial (not exponential) in Walsh weight.

Computation 22 · §18.3 · [NUMERICAL]
Higher-order Dirac commutators.

$\text{ad}_D^k(\chi_S)$, $k = 1, \ldots, 10$. Result: $\|\text{ad}_D^k(\chi_S)\|^{1/k} \to 2\sqrt{D}$ (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.

Computation 23 · §18.3 · [NUMERICAL]

Fréchet smooth L_round closure.

Measures L_F from equation (259) 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 [116] Theorem 5.2 this gives QGH convergence of the Walsh substrate algebras to $C(S^3)$.

Computation 24 · §18.3 · [PROOF]

χ_a–χ_S discrete-side audit.

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.

Computation 25 · §18.3 · [PROOF]

Friedrich S^3 Dirac spectrum.

Eigenvalues $\pm(n + 3/2)/\ell$, multiplicity $2(n + 1)(n + 2)$ ([113]). 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 ([114]). Continuum-side reference for the spinor extension of b_D.

Computation 26 · §18.3 · [PROOF]

Bijective-bridge no-go.

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. Conclusion: no bijective single-mode bridge closes L_comm.

Computation 27 · §18.3 · [PROOF]

Bergman/Toeplitz framework on $H^2_\alpha(B^2)$.

Full operator-theoretic implementation of [119]. 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 $k = 1$, $\|[T_{\bar{z}}, T_z]\|_{\text{op},\alpha} = 1/(\alpha+3)$ *exactly* (machine-precision match against the analytical formula, J-maximiser $(0, 0)$); 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.

Computation 28 · §18.3 · [NUMERICAL]

Explicit substrate–Toeplitz bridge probe (D = 4).

First explicit instantiation of the candidate substrate ↔ 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 $S = \{0, 2\}$). Conclusion: 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.

Computation 29 · §18.3 · [NUMERICAL]

Mosco-limit reweighting via holomorphic Toeplitz subalgebra.

Restricts the bridge image of Computation 28 to the holomorphic Toeplitz subalgebra of $H^2_\alpha(B^2)$, where all operators commute (multiplication by holomorphic functions on H^2_α is just multiplication, no projection needed). Site-to-holomorphic mapping at $D = 4$: sites $0, 1 \to T_{z_0}, T_{z_1}$; sites $2, 3 \to T_{z_0^2}, T_{z_1^2}$. 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 $\alpha = 0$). Conclusion: holomorphic Toeplitz subalgebra is the structurally right target (abelian preservation), but the simplest Clifford analogue $T_{\bar{z}_a}$ 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.

Computation 30 · §18.3 · [NUMERICAL]

Optimal-α closure diagnostic for the holomorphic bridge.

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 analogue 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.

Computation 31 · §18.3 · [NUMERICAL]
Spinor extension on $H^2_\alpha(B^2) \otimes \mathbb{C}^N$ for L_comm.

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 $\mathrm{Cl}(0,4)$ generators tensored with HOLO Toeplitz, (C) mixed bridge with ANTI-HOLO Toeplitz tensored with $\mathrm{Cl}(0,4)$ 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. Conclusion: 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 SU(2) left-invariant vector fields on H^2_α acting as raising/lowering operators on the Wigner-D monomial basis.

Computation 32 · §18.3 · [NUMERICAL]
SU(2) Dirac operator on $H^2_\alpha(B^2) \otimes \mathbb{C}^2$ for L_comm.

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 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 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 $\|[D, b_{\mathrm{holo}}(\chi_S) \otimes I]\|_{\mathrm{op}}$ 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. Conclusion: the 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.

Computation 33 · §18.3 · [NUMERICAL]

Mosco averaging and α rescaling for L_comm closure.

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. Conclusion: this is the FIRST computation demonstrating monotone convergence of the L_comm ratio towards 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.

Computation 34 · §18.3 · [NUMERICAL]

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

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., $k = 3$: $0.86 \to 0.81 \to 0.96$ 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 $\|\text{bridge}([D, \chi_S]) - [D_\alpha, \text{bridge}(\chi_S)]\|_{\text{op}} \to 0$, 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.

Computation 35 · §18.3 · [NUMERICAL]

L_comm homomorphism-gap test (explicit construction).

Tests the genuine L_comm criterion,

$$\|\text{bridge}([D_{\text{sub}}, \chi_S]) - [D_\alpha, \text{bridge}(\chi_S)]\|_{\text{op}}. \tag{283}$$

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. *Negative finding*: 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. *Structural reading*: 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 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.

Computation 36 · §18.3 · [NUMERICAL]

L_comm finite-N matched-truncation convergence at $D = 8$, $N = 20$.

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 ratio $= \|[D_\alpha, \text{bridge}_{\text{final},k} \otimes I]\|_{\text{op}}/(2\sqrt{k})$. *Finding (finite-$N = 20$)*: 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 *finite-$N$* exponential decay with rate $c \approx 0.93$. *Caveat established by Computations 37–38*: the rate 0.93 is specific to the matched truncation $N(D) \sim 2.5\, D$; 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 book-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.

Computation 37 · §18.3 · [NUMERICAL]

Analytical-numerical bridge validation at $D = 4$, $k = 1$.

Sets up the closed-form infinite-N framework: the Leibniz identity [187] $[J_a, T_F] = T_{J_a F}$ for the SU(2) generators acting on holomorphic Toeplitz symbols on $H^2_\alpha(B^2)$. 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 $\|\sigma_1 a + \sigma_2 b + \sigma_3 c\|^2_{\text{op}} = (|a|^2 + |b|^2 + |c|^2) + \sqrt{X^2 + |Y|^2}$ 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. *Validation*: 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-optimisation sequence over ∂B^2 indexed by (D, k).

Computation 38 · §18.3 · [NUMERICAL]

Infinite-N L_comm gap at $k = k_D$: cutoff-transition discontinuity.

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 $D = 4, 6, 8, 10$, by grid-searching sup over the 3-real-parameter unit sphere in $\mathbb{C}^2$. *Findings*: (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 $k_D = 2 \to 3$ (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.

Computation 39 · §18.3 · [NUMERICAL]

Refined symmetric-monomial bridge: closed-form ratio.

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 minimise gap, the residual is 0.02–0.11 across $k = 1$–11.

Computation 40 · §18.3 · [NUMERICAL]

Two-monomial refined bridge: even-k closure.

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.

Computation 41 · §18.3 · [NUMERICAL]

Three-monomial refined bridge: odd-k closure.

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 $= 1.5 \times 10^{-5}$. 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.

Computation 42 · §18.3 · [NUMERICAL]

Refined-bridge closure verification at higher k.

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 $= 3.6 \times 10^{-4}$; $k = 10$: $m = 6$, $\beta = 3.91$, $\alpha = -2.50$, gap $= 5 \times 10^{-5}$; $k = 11$: $m = 6$, $\beta = -0.44$, $\alpha = 5.24$, gap $= 8 \times 10^{-5}$; $k = 12$: $m = 6$, $\beta = -1.29$, $\alpha = -9.52$, 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.

Computation 43 · §14.2 · [PROOF]

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

Sets up the Fano-plane octonion multiplication explicitly and verifies non-associativity $(e_1e_2)e_4 \neq e_1(e_2e_4)$. 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.

Computation 44 · §14.2 · [PROOF]
$\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ multiplication setup.

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, $e_1, \ldots, e_7$) 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).

Computation 45 · §14.2 · [PROOF]
Extended chiral idempotent f on $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$.

Constructs the extended chiral idempotent $f = (I + i_C \otimes I_H \otimes \omega_6)/2$ where $\omega_6 = L_{e_1} \cdots L_{e_6}$ is the Cl(0,6) volume element ($\omega_6^2 = -I$). Verifies $f^2 = f$ exactly, $\mathrm{rank}(f) = 32$ (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 §14.2 to fix the chiral half of each per-generation Hilbert space $\mathcal{H}_k$ at dimension 16.

Computation 46 · §14.2 · [PROOF]
Three quaternionic sub-algebras containing fixed $C \subset \mathbb{O}$ (structural seed for $N_{\mathrm{gen}} = 3$).

Fixes $C = \mathrm{span}\{1, e_1\} \subset \mathbb{O}$ 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 $V_7 \setminus \mathrm{span}\{e_1\} = \mathrm{span}\{e_2, \ldots, e_7\}$ 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 combinatorial structure on $\mathbb{O}$

that PST reads as $N_{\text{gen}} = 3$, consistent with Furey 2014's six-SU(3)-triplet decomposition of $\text{Cl}(0,6)$; colour-vs-generation distinguishability on $\mathbb{O}$ alone is non-trivial ($\text{Stab}_{G_2}(\hat{\tau}) \cong \text{SU}(3)$ acts on $\hat{\tau}^{\perp}$ as the fundamental, mixing the three sub-algebras) and is resolved in the ambient $\mathbb{C}\otimes\mathbb{H}\otimes\mathbb{O}$ (Computation 48).

Computation 47 · §14.1 · [PROOF]
Per-Q_k Cl(0,2) chiral structure.

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 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).

Computation 48 · §14.1 · [PROOF]
Three disjoint generation sectors in $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$.

Resolves the colour-vs-generation distinguishability of $N_{\text{gen}} = 3$ in the Dixon-algebra ambient: 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.

Computation 49 · §14.2 · [PROOF]

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

For the unit direction $\hat{\tau} \in V_7$ selected at modal sublimation, verifies that the cross-product-induced operator $J\colon \hat{\tau}^\perp \to \hat{\tau}^\perp$ 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 $\hat{\tau}$-stabiliser $\mathrm{SU}(3)_{\mathrm{col}} \subset G_2$. Underwrites the colour-SU(3)-as-stabiliser mechanism of §14.2 and the disjoint-generation-sector decomposition of Computation 48.

Computation 50 · §18.3 · [PROOF]

Rigorous single-mode L_round closure.

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

Computation 51 · §18.3 · [PROOF]

Block-decomposition identity and tail-sum reduction.

Diagonalises D_{sub} via spectral projectors $P_\pm = (I \pm D_{\mathrm{sub}}/\sqrt{D})/2$. Verifies to machine precision that $\mathrm{ad}_{D_{\mathrm{sub}}}$ 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 $\|\mathrm{ad}_D^j(T)\|_{\mathrm{op}} = (2\sqrt{D})^j \cdot \max(\|T_{+-}\|, \|T_{-+}\|)$ for every T and every $j \ge 1$. Consequence: $L_F(T) = \max(\|T_{+-}\|, \|T_{-+}\|) \cdot M(D)$ where $M(D) = \sup_j (2\sqrt{D})^j/j!$. For $T_{\mathrm{tail}} = \sum_{|S|>k_D} \chi_S$ at $k_D = 2$, $D \in \{4,5,6,7,8\}$: $\|T_{\mathrm{tail}}\|_{\mathrm{op}} = 2^D - \sum_{j \le k_D} \binom{D}{j}$ and $\max(\|(T_{\mathrm{tail}})_{+-}\|, \|(T_{\mathrm{tail}})_{-+}\|) = 2^{D-1} - D$

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 $P_+T_{\text{tail}}P_-$.

Computation 52 · §18.3 · [PROOF]
Closed-form $\sigma_1 = 2^{D-1} - D$ for the Walsh-tail off-diagonal block.

SVD profile of $P_+T_{\text{tail}}P_-$ across $D \in \{4, \ldots, 10\}$ at cutoff $k_D = 2$. The dominant singular value matches $\sigma_1 = 2^{D-1} - D$ exactly at every D; the dominant right singular vector $v_1 = \frac{1}{\sqrt{2}}\big[|0\rangle - \frac{1}{\sqrt{D}}\sum_a |e_a\rangle\big]$ is the $-\sqrt{D}$-eigenvector of D_{sub} confined to weights $\{0, 1\}$, and the corresponding left singular vector u_1 is its $+\sqrt{D}$ counterpart. Direct calculation in this two-state model yields $\sigma_1 = (\alpha(D) - \beta(D))/2$ with $\alpha(D) = 2^D - 1 - D - \binom{D}{2}$ (the T_{tail} eigenvalue on $|0\rangle$) and $\beta(D) = -\binom{D-1}{2}$ (on each $|e_a\rangle$). Simplification gives $\sigma_1 = 2^{D-1} - D$ *exactly*, closing Lemma 2-prime on the S_D-trivial component of $P_+T_{\text{tail}}P_-$. The second singular value $\sigma_2 = D - 2$ (numerical here; analytically closed in Computations 54, 55) is bounded by an exponential gap from σ_1.

Computation 53 · §18.3 · [PROOF]
Algebraic identity $P_+T_{\text{tail}}P_- = (1/(2\sqrt{D}))\, P_+ [D_{\text{sub}}, T_{\text{tail}}]$.

Verifies to machine precision the algebraic identity $P_+T_{\text{tail}}P_- = (1/(2\sqrt{D}))\, P_+ [D_{\text{sub}}, T_{\text{tail}}]$, obtained by substituting $D_{\text{sub}}T_{\text{tail}}D_{\text{sub}} = D \cdot T_{\text{tail}} - D_{\text{sub}}[D_{\text{sub}}, T_{\text{tail}}]$ (which follows from $D_{\text{sub}}^2 = D \cdot I$) into the expansion of $P_+T_{\text{tail}}P_-$. Closed form on the weight-1 standard-rep input alone: $\sigma_{\max} = (D-2)/\sqrt{2} = |\beta(D) - \gamma(D)|/(2\sqrt{2})$ with $\gamma(D) = c(2) = (D-2)(5-D)/2$ the T_{tail} eigenvalue on weight-2 states, achieved via the exact cancellation $D_{\text{sub}}w_2 = D \cdot v$ for $w_2 = \sum_{a<b}(c_b - c_a)|e_a + e_b\rangle$ when $\sum c_a = 0$. The representation that commutes with D_{sub} is the

fermionic-signed one; under it $\sigma_2 = D - 2$ exactly (Computations 54, 55).

Computation 54 · §18.3 · [NUMERICAL]
Identification of the S_D-irrep that saturates $\sigma_2 = D - 2$.

Implements the character-projector decomposition of T^{+-}_{tail} under the fermionic-signed S_D representation $\rho_F(g)|S\rangle = \text{sgn}(g \mid S)|g(S)\rangle$ (the naive bit-permutation rep does *not* commute with D_{sub} because of the Jordan–Wigner σ_z strings; verified $\|[\rho_{\text{bit}}(g), D_{\text{sub}}]\|_{\text{op}} \approx 4.0$ at $D = 4$ versus $\|[\rho_F(g), D_{\text{sub}}]\|_{\text{op}} = 0$ to machine precision). Builds $P_\lambda = (\dim\lambda/|S_D|)\sum_g \chi_\lambda(g)\,\rho_F(g)$ for each two-row irrep $[D - j, j]$ with $j = 0, 1, \ldots, \lfloor D/2 \rfloor$, then computes $\sigma_{\max}$ within each isotypic block. *Findings:* the standard rep $[D - 1, 1]$ saturates $\sigma_{\max} = D - 2$ exactly at $D = 4, 5, 6, 7$; every higher two-row irrep gives $\sigma_{\max} = 0$ identically. The multiplicity of $[D - 1, 1]$ under ρ_F is 2 (Pieri on $\Lambda^w(\mathbb{C}^D)$ assigns it to H_1 and H_2), so the multiplicity-space matrix M is 2×2; numerically $\text{rank}(M) = 1$ with singular values $(D - 2, 0)$.

Computation 55 · §18.3 · [PROOF]
Closed-form $\sigma_2 = D - 2$ on $[D - 1, 1]$, completing Lemma 2-prime.

Exhibits the explicit S_{D-1}-fixed vectors $u_1 = |e_0\rangle - (1/D)\sum_a |e_a\rangle \in [D - 1, 1] \cap H_1$ and $u_2 = -\sum_{b>0} |\{0, b\}\rangle \in [D - 1, 1] \cap H_2$ and verifies to machine precision (for $D = 4, \ldots, 8$) the algebraic identities $D_{\text{sub}}u_1 = u_2$, $D_{\text{sub}}u_2 = Du_1$ (no weight-3 leakage; the JW orderings on each triple $\{0, b, c\}$ cancel), $T_{\text{tail}}u_1 = \beta\, u_1$ with $\beta = -\binom{D-1}{2}$, $T_{\text{tail}}u_2 = \gamma\, u_2$ with $\gamma = (D - 2)(5 - D)/2$. Applying $T^{+-}_{\text{tail}} = (1/(2\sqrt{D}))\, P_+\, [D_{\text{sub}}, T_{\text{tail}}]$ collapses the multiplicity-space matrix in the orthonormal basis $e_1 = u_1/\|u_1\|$, $e_2 = u_2/\|u_2\|$ (with $\|u_1\|^2 = (D - 1)/D$, $\|u_2\|^2 = D - 1$) to $M = \frac{D-2}{2}\binom{1}{1}\begin{pmatrix}-1 & 1\end{pmatrix} = \frac{D-2}{2}\begin{pmatrix}-1 & +1\\ -1 & +1\end{pmatrix}$, rank 1 with singular values $(D - 2, 0)$. The saturating right singular vector

is $v_\star = (1/\sqrt{2})(-e_1 + e_2)$, the left singular vector is $u_\star = (1/\sqrt{2})(e_1 + e_2)$. Combined with Computation 52 this gives Lemma 2-prime as a theorem: $\|T^{+-}_{\text{tail}}\|_{\text{op}} = 2^{D-1} - D$ *exactly* for every $D \geq 4$ at $k_D = 2$.

Computation 56 · §18.3 · [PROOF]

Closed-form L_comm ratio for the single-monomial bridge (Lemma 5(a)).

Derives in closed form the L_comm ratio for the single-monomial bridge symbol $f(w) = w^m$ on the holomorphic Toeplitz algebra of the weighted Bergman space ($w = z_0 z_1$). Since $J_z f = 0$ (f is symmetric in z_0, z_1) and $J_\pm f$ are antiholomorphic mirror pairs, M_f is anti-diagonal in the Pauli basis with $\|M_f\|_{\text{op}} = \max(|z_0|^2, |z_1|^2) \cdot |f'(w)|$. The 1D optimisation $r^2 = (m+1)/(2m)$ on ∂B^2 gives $\text{ratio}(m) = \sup \|M_f\|_{\text{op}} / \sup |f| = m^{1-m}\,(m+1)^{(m+1)/2}\,(m-1)^{(m-1)/2}$, verified to machine precision for $m = 1, \ldots, 10$. At $m = 1$ the formula evaluates to $2 = 2\sqrt{1}$, so the single-monomial bridge achieves exact L_comm closure at weight $k = 1$ with zero refinement parameters. For $m \geq 2$ the ratio differs from $2\sqrt{m}$ (undershoot at $m = 2$, overshoot at $m \geq 3$; asymptotically $\text{ratio}(m) \sim m + 1/2$).

Computation 57 · §18.3 · [PROOF]

Lemma 5(b): existence of closing parameters for every k via IVT.

Establishes the existence side of the refined-bridge closure for the 2-monomial family $f_\alpha(w) = w^m + \alpha\, w^{m+1}$. The ratio $R_m(\alpha) := \sup \|M_{f_\alpha}\|_{\text{op}} / \sup |f_\alpha|$ is continuous in α with $R_m(0) = \text{ratio}(m)$ (Computation 56) and $\lim_{|\alpha|\to\infty} R_m(\alpha) = \text{ratio}(m+1)$. Since $\text{ratio}(m)$ is strictly increasing in m, the substrate target $2\sqrt{k}$ is bracketed by $\text{ratio}(m(k)) \leq 2\sqrt{k} \leq \text{ratio}(m(k)+1)$ for a unique $m(k) \in \mathbb{N}$; Bolzano then gives $\alpha(k) \in \mathbb{R}$ with $R_{m(k)}(\alpha(k)) = 2\sqrt{k}$ exactly. Verifies the bracket condition for $k = 1, \ldots, 12$ (extends to all k

via ratio$(m) \sim m + 1/2$) and computes the closing $\alpha(k)$ by IVT bisection: $\alpha(2) \approx 0.74$, $\alpha(3) \approx 20.7$, $\alpha(4) \approx 1.74$, …, $\alpha(12) \approx 1.38$, all with R-target gap at machine precision. *Theorem (Lemma 5(b))*: for every $k \geq 1$, there exist $m(k) \in \mathbb{N}$ and $\alpha(k) \in \mathbb{R}$ such that the 2-monomial bridge symbol $f_k(w) = w^{m(k)} + \alpha(k)\, w^{m(k)+1}$ achieves exact L_comm closure ratio $2\sqrt{k}$.

Computation 58 · §18.3 · [PROOF]

Lemma 5(c): closed-form parametric expression for $\alpha(k)$.

Closes Lemma 5 with a closed-form parametric expression for the L_comm closing parameter. The substitution $q := 2r^2 - 1$ parameterises the critical point on ∂B^2, and the critical-point equation gives $\alpha(q) = 2m(1 - mq)/[(m + 1)\sqrt{1 - q^2}((m + 1)q - 1)]$ for $q \in (1/(m + 1), 1/m)$. Substituting back, the closing equation $R_{m(k)}(q^*) = 2\sqrt{k}$ is an algebraic equation in q^* of degree $2m + 3$ after rationalisation, so $q^*(k)$ is an algebraic number and $\alpha(k) = \alpha(q^*(k))$ is a rational function of it. Verified for $k = 2, \ldots, 12$: the parametric expression matches Computation 57's IVT-bisection $\alpha(k)$ values to six decimal places. The endpoint behaviour $q^* \to 1/m$ gives the single-monomial limit, $q^* \to 1/(m + 1)$ the dominant-$(m + 1)$ limit. *Lemma 5 (closure summary)*: parts (a) (single-monomial closed form, Computation 56), (b) (existence, Computation 57), and (c) (parametric closed form, Computation 58) jointly close the L_comm refined-bridge ratio-matching step as a complete analytical theorem. The remaining open work for the full L_comm closure is structural: Berezin-transform compatibility, spectral-action invariance, and the uniform closure rate as $D \to \infty$.

Computation 59 · §18.3 · [NUMERICAL]

Lemma 7: uniform-in-k closure of the refined bridge at finite N (numerical).

Numerical investigation of the truncation gap of the exact 2-monomial bridge (with $\alpha(k)$ from Lemma 5(c)) on the truncated weighted Bergman space $H^2_\alpha(B^2)^{(N)}$, as a function of N and weight k. Builds the SU(2) Dirac $D_\alpha = J_a \otimes \sigma_a$ in the orthonormal monomial basis, applies the bridge T_{f_k}, and measures $\mathrm{gap}(k, N) := \|[D_\alpha, T_{f_k} \otimes I]\|_{\mathrm{op}} / \|T_{f_k}\|_{\mathrm{op}} - 2\sqrt{k}$. Findings: at $k = 1$ the gap is essentially exactly $-1/N$ (least-squares fit gives $c_1 = -0.9968$, $c_2 = -0.2985$ with max residual 1.7×10^{-4}; the closed form derived in Computation 60 is $\mathrm{gap}(1, N) = 2\sqrt{1 - 1/N} - 2$ exactly); at $k \geq 2$ the gap is also $O(1/N)$ but with k-dependent constants (leading c_1 values -0.43, $+1.15$, $+0.70$ at $k = 2, 3, 4$) and a less clean two-term fit. At the matched scaling $N(D) \sim 2k_D \sim 2\sqrt{D}$ this implies $\gamma_D = O(1/\sqrt{D})$ at $k = 1$, polynomial L_comm closure rate, weaker than the exponential L_round rate but uniformly controlled. The full uniform-in-k analytical bound for $k \geq 2$ requires careful expansion of T_{f_k} on the "boundary" modes (total degree close to N) of the truncated Bergman space.

Computation 60 · §18.3 · [PROOF]

Lemma 7(a): closed-form $\mathrm{gap}(1, N) = 2\sqrt{1 - 1/N} - 2$ for the bare bridge.

Promotes Computation 59's $k = 1$ numerical fit to an exact closed-form theorem. For canonical Bergman weight $\alpha = 0$ on B^2 with orthonormal monomial basis $e_{a,b}$ restricted to $a + b \leq N$, the holomorphic Toeplitz operators $T_{z_0^p z_1^q}$ act as diagonal shifts $e_{a,b} \to (\text{coefficient}) \cdot e_{a+p,b+q}$ (orthogonal target). Operator norms read off as maximum matrix elements: $\|T_w\|_{\mathrm{op}} = N/[2\sqrt{(N+1)(N+2)}]$ (maximiser $a = b = (N-2)/2$ for even N) and

$\|T_{z_0^2}\|_{\mathrm{op}} = \sqrt{N(N-1)/((N+1)(N+2))}$ (maximiser $a = N-2$, $b = 0$). The commutator $[D_\alpha, T_w \otimes I_2] = T_{z_0^2} \otimes \sigma_+ + T_{z_1^2} \otimes \sigma_-$ is anti-diagonal in the Pauli basis, giving the closed form $R(N) = 2\sqrt{1-1/N}$ exactly for even N. *Theorem (Lemma 7(a))*: $\mathrm{gap}(1, N) = 2\sqrt{1-1/N} - 2$ exactly, with Taylor expansion $-1/N - 1/(4N^2) - 1/(8N^3) - 5/(64N^4) - O(1/N^5)$. Verified against Computation 59's numerical SU(2)-Dirac measurement to ~ 6 decimal places at every $N \in \{6, 8, \ldots, 22\}$. Closure rate at matched scaling $N(D) \sim 2\sqrt{D}$: $\gamma_D(k=1) = -1/(2\sqrt{D}) + O(1/D)$, polynomial in D, weaker than the exponential L_round rate but uniformly controlled.

Computation 61 · §18.3 · [PROOF]

Lemma 6 (Berezin compatibility) closed: $B(T_f T_g^*)(z) = f(z)\,\overline{g(z)}$ exactly.

Closes Lemma 6 (Berezin-transform compatibility) for the refined bridge. *Theorem (Lemma 6)*: for holomorphic symbols f, g on B^2 and the canonical weighted Bergman space $H^2_\alpha(B^2)$, $B(T_f T_g^*)(z) = f(z)\,\overline{g(z)}$ EXACTLY for every z in the open ball. Proof: the reproducing-kernel adjoint identity $T_g^* k_z = \overline{g(z)}\, k_z$ follows from $\langle h, T_g^* k_z\rangle = \langle T_g h, k_z\rangle = (T_g h)(z) = g(z)h(z) = \langle h, \overline{g(z)}\, k_z\rangle$ for all $h \in H^2_\alpha$. Substituting, $B(T_f T_g^*)(z) = \langle k_z, T_f T_g^* k_z\rangle / \|k_z\|^2 = \overline{g(z)}\, B(T_f)(z) = \overline{g(z)}\, f(z)$, using the standard $B(T_f)(z) = f(z)$ for holomorphic f. Applied to the refined bridge $b_C(\chi_S) = T_{f_{|S|}}$ with f_k holomorphic polynomial in $w = z_0 z_1$: $B(b_C(\chi_S)\, b_C(\chi_T)^*)(z) = f_{|S|}(z)\,\overline{f_{|T|}(z)}$ EXACTLY. On the truncated Bergman $H^2_\alpha(B^2)^{(N)}$ the identity holds up to an $O(|z|^{2N})$ boundary correction that decays exponentially in N for z in the open ball. Verified numerically at every tested $z \in B^2$ (interior) and every $k, l \in \{1, 2, 3\}$: gap below numerical precision at $N = 18$. The bridge therefore realises the substrate's weight-class structure on the Bergman side as a holomorphic-symbol algebra, and the Berezin transform recovers the symbol algebra exactly.

Computation 62 · §18.3 · [PROOF]

Lemma 8 (spectral-action invariance corollary) at matched scaling.

Closes Lemma 8 (spectral-action invariance under the refined bridge b_C) as a corollary of Lemmas 5, 6, 7(a). Substrate Dirac D_{sub} has eigenvalues $\pm\sqrt{D}$ each with multiplicity 2^{D-1}; Bergman SU(2) Dirac D_α has Friedrich spectrum $\pm(n+3/2)/\ell$ with multiplicity $2(n+1)(n+2)$. Matched scaling: $\Lambda = \sqrt{D}$, $\ell = (3\cdot 2^D)^{1/3}/\Lambda$ so that the Bergman level count up to Λ (without multiplicity) equals 2^D = substrate Hilbert-space dimension, the count with the multiplicity $2(n+1)(n+2)$ for both signs being $4\cdot 2^D$, a D-independent factor. *Theorem (Lemma 8)*: conditional on Lemmas 5, 6, 7(a), the Connes spectral action $S(D,\Lambda,f) = \operatorname{Tr} f(D/\Lambda)$ of the refined-bridge image is commensurate with the substrate spectral action at the matched scaling, up to a cutoff-function redefinition. Both sides scale as 2^D and their ratio is bounded independently of D. Verified numerically with Gaussian cutoff $f(x) = e^{-x^2}$ for $D \in \{6, 8, 10, 12\}$: $S_{\text{sub}}/2^D = e^{-1}$ exactly; $S_{\text{Bergman}}/2^D \to$ constant ≈ 21.7 (independent of D) at $j = 1$, ≈ 54.2 at $j = 2$, ≈ 189.7 at $j = 3$. Standard QGH-convergence statement applied to spectral actions. Together with Lemma 7(b)'s direct-gap closure, all six L_comm sub-lemmas are closed; that closure lands in Computation 65, after the fit-stability diagnostics of Computations 63 and 64 establish that the $c_1(k)$-extraction route is unreliable and must be replaced by the direct-gap bound.

Computation 63 · §18.3 · [NUMERICAL]

Lemma 7(b): leading $c_1(k)$ for $k = 1, \ldots, 12$ (2-parameter least-squares fit).

Numerical extraction of the leading constant $c_1(k)$ in $\text{gap}(k, N) = c_1(k)/N + c_2(k)/N^2 + O(1/N^3)$ via a two-parameter least-squares fit on N adapted to the bridge degree. Confirms Lemma 7(a)'s $c_1(1) = -1$ exactly with fit residual

$\sim 10^{-5}$. Reports $|c_1(k)| \leq 2.7$ across $k = 1, \ldots, 12$. The fit stability of this extraction is interrogated in Computation 64; the foundational uniform-in-k gap bound is established directly (without c_1 extraction) in Computation 65.

Computation 64 · §18.3 · [NUMERICAL]

Lemma 7(b) numerical-fit-stability diagnostic.

Re-extracts $c_1(k)$ on the same N range used by Computation 63 (capped at $N \leq 34$) under a three-parameter least-squares fit $\text{gap}(k, N) \approx c_1/N + c_2/N^2 + c_3/N^3$. For $k = 1$ the two- and three-parameter fits agree at $c_1 = -1.000$, matching the closed-form theorem of Computation 60. For $k \geq 2$ the two fits DISAGREE: factor 2–10 in magnitude and sign-flip in 9 of 23 measured k values. Diagnostic conclusion: the per-k leading constant $c_1(k)$ for $k \geq 2$ cannot be reliably extracted from $N \leq 34$ by fit-based methods; the leading $1/N$ term is not yet separated from higher-order structure when N is comparable to the bridge degree $m(k) + 1$. Hence the foundational uniform-in-k bound (Lemma 7(b)) cannot rest on $c_1(k)$ extraction at moderate N; it must be established directly (Computation 65).

Computation 65 · §18.3 · [NUMERICAL]

Lemma 7(b) direct-gap closure: $|\text{gap}(k, N)| \cdot N$ bounded.

Bypasses the ill-conditioned $c_1(k)$ extraction by directly measuring the gap-times-N quantity $G(k, N) := |\text{gap}(k, N)| \cdot N$ across $k \in [1, 16]$ and $N \in [12, 32]$ (159 measured pairs). Finding: $\max_{(k,N)} G(k, N) = 5.249$ attained at $(k, N) = (15, 20)$; median $G = 1.029$. *Scope corrected by Computation 150*: this window includes points with $k > k_D$, and the reported maximum at $(15, 20)$ is one of them ($k_D = 10$ there), so 5.249 is not a bound on anything the theory uses. Restricted to the operative regime $k \leq k_D$ the sharp constant is 3.7; the bound is *not* uniform over all (k, N), Computation 150 measuring 22.0 at

$(39, 22)$, a point above k_D. With that restriction the matched scaling gives the L_comm closure rate $\gamma_D \leq 3.7/(2\sqrt{D}) = O(1/\sqrt{D})$ uniformly in $k \leq k_D$. Structural backing: both $\|T_{f_k}\|_N^2$ and $\|[D_\alpha, T_{f_k} \otimes I_2]\|_N^2$ approach their bulk values at rate $O(1/N)$ by the classical Bergman–Toeplitz convergence theorem for polynomial symbols (Engliš, Zhu), and their ratio inherits the rate with bounded constant. This closes Lemma 7(b) at the foundational level: the bound uniform in $k \leq k_D$ holds directly without per-k extraction, and Rieffel's QGH-convergence theorem receives the rate it requires.

Computation 66 · §13.7 · [PROOF]

$\mathrm{Cl}(1,3) \cong M_2(\mathbb{H})$ and the internal $\mathrm{SU}(2)$ from right-$\mathbb{H}$.

Structural verification that the substrate's emergent 4-d Lorentzian spacetime $M = \mathbb{R} \times S^3$ (derived from P1–P3 via Mosco convergence, Computation 4) carries an internal SU(2) on its Dirac spinor sector. Works in $\mathbb{H}^2 \cong \mathbb{R}^8$ (the $M_2(\mathbb{H})$ module) to avoid antilinearity confusion of the $\mathbb{C}^4$ chiral basis. Builds $\mathrm{Cl}(1,3)$ γ-matrices in the $M_2(\mathbb{H})$ representation (16 anticommutators match $\mathrm{diag}(+1,-1,-1,-1)$ to machine precision); builds right-$\mathbb{H}$ operators R_{iH}, R_{jH}, R_{kH} on $\mathbb{H}^2$; verifies right-$\mathbb{H}$ commutes with $\mathrm{Cl}(1,3)$ (the centraliser structure that makes SU(2) internal, [146]); verifies quaternion algebra $R_q^2 = -I$, $R_j R_i = R_k$; shows $\{R_q/2\}$ generates $\mathfrak{su}(2)$ with $[R_{iH}, R_{jH}] = -2R_{kH}$; verifies the Lorentz pseudoscalar commutes with R_q so the SU(2) is chirality-preserving. This establishes the SPACETIME-side spinor SU(2) structure derived from P1–P3; the identification of this SU(2) with the $\mathrm{SU}(2)_L$ factor of $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ on the internal F-side is via Dixon's hyperspinor framework $T = \mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$, an established structural synthesis beyond $\mathrm{Cl}(0,6)$ alone (cf. [146], [121] 'A bigger picture' §).

Computation 67 · §12.2 · [PROOF]

One-loop Higgs cancellation: scope check of eq. (129)**.**

Combinatorial verification of the PST scalar-sector cancellation argument (§12.2, equations (128) and (129)). Confirms the two exact binomial identities $\sum_{k\,\text{even}} \binom{D}{k} = \sum_{k\,\text{odd}} \binom{D}{k} = 2^{D-1}$ across $D \in \{4, 6, 8, 10, 12, 14, 16\}$. Including the zero mode in the bosonic count gives $N_B - N_F = 0$ exactly. *Excluding* the zero mode would instead give $N_B^{(\text{nonzero})} - N_F = -1$ for every D, contributing a residual $-M_*^2/(16\pi^2)$ to eq. (129) rather than 0. Added to the self-loop $+(3/2)M_*^2/(16\pi^2)$ of eq. (128), this gives total $\delta m_h^2 = M_*^2/(32\pi^2)$, conflicting with the predicted $m_h = M_*\sqrt{6}/(8\pi)$. The off-by-one is resolved by the substrate-vs-emergent layer distinction the body now adopts (Computations 93, 97): the substrate Walsh zero mode $\chi_\emptyset$ maps under Π to the constant function on M (the vacuum direction), and is structurally distinct from the emergent 4-component SM Higgs doublet generated by inner fluctuations on $M \times F$. Counting the substrate zero mode in N_B^{full} therefore does not double-count the emergent doublet, which is carried separately by $C_H = 6\lambda = 3/2$ in eq. (128). With $N_B^{\text{full}} - N_F = 0$ from the substrate trace and the $+(3/2)M_*^2/(16\pi^2)$ self-loop supplying the Higgs mass, total $\delta m_h^2 = (3/2)M_*^2/(16\pi^2)$, recovering $m_h = M_*\sqrt{6}/(8\pi)$ and confirming $M_* \approx 1.285$ TeV as a parameter-free one-loop derivation. The three mechanisms once entertained here, (a) the Higgs doublet's three eaten Goldstones at the EW scale, (b) reinterpreting the zero mode as part of the cancellation sum, and (c) a complementation/parity pairing of the non-zero Boolean modes, are superseded by this resolution and are walked through and set aside in Computation 69.

Computation 68 · Chapter 15 · [NUMERICAL]

Casimir-coefficient ξ kernel-shape sensitivity.

Tests the sensitivity of the parameter-free Casimir-correction coefficient $\xi = (15/\pi^2) \cdot M_4/(d_0^2 M_2)$ (eq. (213)) to the assumed projection-kernel shape, by evaluating the second and fourth moments M_2, M_4 for six kernel families: Gaussian (book canonical), squared Lorentzian, single-sided exponential, sech-squared, compact parabolic, and compact quartic-bump. Findings: the d^{-6} power-law exponent is preserved across every kernel (categorical prediction, kernel-independent). The coefficient ξ varies from 2.0 (compact quartic-bump) to 21.0 (sech^2) across kernels with exponential or compact-support tails (factor ~ 10 range); the squared Lorentzian gives $\xi = 378.7$ from its $1/x^4$ heavy tail (physically implausible for a discreteness scale). The diagnostic d_0 bound from nanometre-Casimir data scales as $1/\sqrt{\xi}$, giving the bound range $d_0 \lesssim 1.5$ nm (sech^2) to $d_0 \lesssim 5$ nm (compact quartic) at the current 0.19% data, with the Gaussian value 2.3 nm sitting near the middle. Order of magnitude (few nm) is robust; precise numerical bound is kernel-conditional with factor-of-~ 3 spread. Honest reading: "$d_0 \lesssim$ 1.5–5 nm depending on the projection-kernel choice; the Gaussian (canonical) choice gives 2.3 nm." The derived value $d_0 = 2.2$ am (Computation 201) lies nine orders below every bound in the range.

Computation 69 · §12.2 · [PROOF]

Higgs cancellation: detailed analysis of the three Computation 67 resolutions.

Carries out the explicit walkthrough of all three resolutions listed in Computation 67 for the off-by-one in eq. (129). (a) Higgs-doublet eaten Goldstones: the resolution is at the EW scale (broken phase), while the cancellation argument is at $M_* \gg v$ (symmetric phase), dimensionally a category error; not

constructible without an explicit RG-matching argument the book does not provide. (b) Zero mode included in the cancellation sum: gives $N_B - N_F = 0$ exactly, but no leftover $+M_*^2/(16\pi^2)$ to *be* the Higgs self-mass, total $\delta m_h^2 = 0$, $m_h = 0$ at one loop; a no-op resolution. (c) Complementation $S \to S^c$ pairing: for even D, $|S|$ and $|S^c|$ have the same parity so complementation does *not* cross-pair bosons with fermions; for odd D the pairing exists structurally but the zero-mode / full-set pair is broken if the zero mode is set aside as the Higgs/order-parameter direction. None of these three resolutions, as constructible from current PST postulates without ad hoc structure, closes the off-by-one. The fourth resolution does: the substrate Walsh zero mode $\chi_\emptyset$ maps under Π to the constant function on M (vacuum direction), structurally distinct from the emergent 4-component SM Higgs doublet generated by inner fluctuations on $M \times F$ (Computations 93, 97). Including the zero mode in N_B does not double-count $C_H = 6\lambda = 3/2$, and $M_* = 4\pi m_h \sqrt{2/3}$ stands as a parameter-free one-loop derivation.

Computation 70 · §18.1.6 · [SURROGATE]

K-fixed Berry–Esseen rate for A7.

Verifies the Berry–Esseen-type convergence rate underlying the A7 reduction lemma: the empirical Boolean Dirichlet form converges to the continuum L^2-gradient at CLT slope ≈ -0.467 on a single compact-support test family in $K = [-2, 2]^3 \subset \mathbb{R}^3$. Three test functions (Gaussian, C^∞ bump, polynomial-cosine). Establishes the K-fixed Berry–Esseen rate; Computation 83 extends to K-uniform. Stands in for: the emergence-determined ρ_n, by i.i.d. uniform sites with the identity embedding.

Computation 71 · §12.2.2 · [NUMERICAL]

$Z^2 \approx e^{-1}$ candidate identification via substrate spectral action.

Tests whether the one-loop near-coincidence $Z^2 := \lambda_{\mathrm{SM}}(M_*)/b(M_*) \approx e^{-1}$ (0.8% deviation) lifts to a structural identification via the substrate spectral-action ratio $S_{\mathrm{sub}}/2^D = e^{-1}$ (exact, D-independent at matched scaling $\Lambda = \sqrt{D}$). Verifies the substrate side exactly for $D \in \{4, 6, 8, 10, 12, 14, 16\}$; computes one-loop SM RGE running $\lambda_{\mathrm{SM}}(v) \to \lambda_{\mathrm{SM}}(M_*)$ across $\{\lambda, y_t, g_1, g_2, g_s\}$; reports $\lambda_{\mathrm{SM}}(M_*) \approx 0.087$. Identifies the spectral-action Higgs-quartic identity as the open structural conjecture. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 72 · §12.2.2 · [NUMERICAL]

Higgs perturbation and ϕ^4 coefficient at matched scaling.

Carries out the symbolic Taylor expansion of the substrate spectral action under a Higgs-like perturbation $D_\phi = D_{\mathrm{sub}} + y\,\phi\,\Sigma_X$ at matched scaling, extracts the ϕ^4 coefficient $e^{-1} \cdot y^4/(2D^2)$, and reduces the Z^2 identification to the single-scalar condition $y/\Lambda = 2^{-1/4} = \mu_{\mathrm{site}}^{1/4}$ (the fourth root of the per-site Bernoulli probability of P1). Identifies the open structural content: deriving $y/\Lambda = 2^{-1/4}$ from CC inner fluctuation.

Computation 73 · Chapter 15 · [NUMERICAL]
Casimir ξ from the substrate Boolean kernel.

Sharpens Computation 68's kernel-shape sensitivity by deriving $\xi = (15/\pi^2)M_4/(d_0^2 M_2)$ from the substrate's actual projection kernel (the binomial Boolean kernel) at finite $|D|$. Computes the moments M_2, M_4 of the Boolean kernel as binomial-sum identities; identifies the leading non-Gaussian correction; extrapolates to the matched-scaling limit. Confirms the Gaussian kernel as the natural matched-scaling limit; documents finite-$|D|$ corrections at the $O(1/D)$ level.

Computation 74 · §12.2.2 · [NUMERICAL, superseded by Computation 100]
First-cut CC inner-fluctuation for $Z^2 = e^{-1}$.

Tests the simplest natural single-generator CC inner fluctuation $A = i\phi[D_{\text{sub}}, \omega]/2$ with $\omega = e_1 e_2$ (quaternion volume element). Computes the perturbed Dirac $D_\phi^2 = (D + 2\phi^2)I$ and the ϕ^4 coefficient $2e^{-1}/D^2$. Matching Computation 72 target gives $y = \sqrt{2}$ (D-independent) hence $y/\Lambda = \sqrt{2/D}$, which DECREASES with D. Target is $2^{-1/4}$ (D-independent). Conclusion: natural single-generator inner fluctuation does *not* close $Z^2 = e^{-1}$. Identifies three possible resolutions (numerical coincidence, multi-generator A, alternative $\Lambda(D)$).

Computation 75 · §16.7 · [NUMERICAL]
Rate integral j_τ as Bernoulli large-deviation.

Sharpens the rate-integral closure (book eq. (229)) by identifying j_τ as a standard large-deviation integral under the Bernoulli measure. $j_\tau \sim \exp(-|D| \cdot I(\tau/\sqrt{|D|}))$ with I the Cramér rate function. Computes saddle-

point asymptotic for a representative $T(C)$ form; documents the rate integral's structural form. Closes the rate-problem framework for downstream applications (cosmological constant, Computation 82).

Computation 76 · §14.3 · [PROOF]

Postulate-level V_7-direction constraint as a Yukawa-hierarchy source.

Tests whether constraining the direction of $v(a) \in V_7$ to the three quaternionic sub-algebras Q_k of $\mathbb{O}$ containing $\hat{\tau}$ lifts the structural-scope theorem. Result: maximal G_2-symmetry preserves S_3-symmetry across the three Q_k, so this single-direction constraint alone does not lift the structural-scope theorem. The threshold-induced dynamical mechanism is taken up in Computation 77.

Computation 77 · §14.3 · [NUMERICAL]

Dynamical Yukawa hierarchy via the directed modal threshold per Q_k.

The directed P3 threshold-crossing rate per quaternionic sub-algebra Q_k gives a hierarchy of generation Yukawa couplings via Boltzmann weighting $\exp(-\beta E_k)$. Reproduces the SM up-type Yukawa hierarchy with $\beta \approx 28.24$ and $(p_1, p_2, p_3) \approx (0.52, 0.35, 0.13)$. Identifies the two inputs requiring structural derivation: β and the (p_k) pattern; structural candidates for each are taken up in Computations 78, 79, 80.

Computation 78 · §14.3 · [NUMERICAL]

Structural candidates for β in the Boltzmann mechanism.

Tests structural candidates for the Boltzmann inverse-temperature β from Computation 77's fit. Candidates: $\dim(\mathrm{SO}(8)) = 28$, $\dim(G_2) = 14$, $4\dim(V_7) =$

28, $\dim(\mathrm{Spin}(7)) = 21$. Numerical match: $\dim(\mathrm{SO}(8)) = 28$ at 0.85% from Computation 77 best-fit 28.24. Motivation: Spin(8) is the rotation group of $\mathbb{O} = \mathbb{R}^8$, with triality permuting three 8-dim reps and three PST generations.

Computation 79 · §14.3 · [NUMERICAL]
$\beta = 28.25 = \dim(\mathrm{SO}(8)) + b_{\mathrm{PST}}$ and $p_2/p_3 = e$.

Sharpens Computation 78's $\beta \approx 28$ to $\beta = \dim(\mathrm{SO}(8)) + b_{\mathrm{PST}} = 28 + 1/4 = 28.25$, matching Computation 77 best-fit to 0.04%. Identifies $p_2/p_3 = e$ in Computation 77's (p_2, p_3) ratio as structurally suggestive of the same e^{-1} that drives Z^2. Locks in $\beta = 28.25$ structurally; p_1 remains as the one residual free parameter.

Computation 80 · §14.3 · [PROOF]
p_1 closure via $\mathrm{Spin}(8) \to \mathrm{SU}(3)_c$ symmetry chain.

Examines whether p_1 in the dynamical mechanism of Computations 77 and 79 can be pinned by the substrate symmetry chain $\mathrm{Spin}(8) \supset \mathrm{Spin}(7) \supset G_2 \supset \mathrm{SU}(3)_c$. Result: $\mathrm{SU}(3)_c$ stabiliser structure forces p_1 to a $T(C)$-contingent value rather than a structural constant. Reclassifies the Yukawa hierarchy as configuration-content (consistent with the structural-scope theorem of §14.3).

Computation 81 · §12.2.2 · [NUMERICAL, superseded by the partition-function route, Computations 87–92/100]
Full CC inner fluctuation $D + A + \varepsilon' JAJ^{-1}$ for Z^2.

Extends Computation 74's single-generator inner fluctuation to include the JAJ^{-1} reality term with KO-2 sign $\varepsilon' = +1$, and tests multi-component A. Result: $y/\Lambda = \sqrt{8/D}$ (vs Computation 74's $\sqrt{2/D}$, factor 2 from reality

doubling), still D-dependent with wrong scaling. Confirms Computation 74's verdict that single- and few-generator inner fluctuations do not close $Z^2 = e^{-1}$. Identifies the structural obstacle: the discrete spectrum's collapse-to-cutoff behaviour.

Computation 82 · §16.5 · [NUMERICAL]

Λ magnitude: order-of-magnitude reclassification as $T(C)$-contingent.

Order-of-magnitude analysis for the cosmological-constant magnitude using Computation 75's Cramér framework. $\rho_{\text{inst}} = j_\tau \cdot \Delta\mathcal{F} \cdot (|D|/V_M)$ with matched-scaling $|D|/V_M \sim 1/d_0^4$. Target j_τ for $\rho_{\text{inst}} = \rho_{\Lambda,\text{obs}}$ comes out to $\sim 10^{-37}$ per substrate configuration at $d_0 = 2.3$ nm (Computation 188 reproduces it; at the derived 2.2 am the matched scaling sends it to $\sim 10^{-73}$, Computation 201); the required rate-function value $I(\alpha) \sim 10^{-110}$ is tiny. *Conclusion*: Λ magnitude depends on the explicit $T(C)$ form plus threshold-scaling $\tau(D)$, both $T(C)$-contingent. Reclassifies Λ magnitude from open derivation to structural contingent (joins Yukawa hierarchy, CKM, d_0). Sign + EoS $w = -1$ remain structural theorems.

Computation 83 · §18.1.6 · [SURROGATE]

K-uniform Berry–Esseen rate verification for A7.

Extends Computation 70's K-fixed verification to K-uniform: seven test-function families with varying support locations, radii, and shapes (Gaussian, C^∞ bump, polynomial-cosine). Mean slope -0.458 ± 0.085 across families; 5/7 families within fitting tolerance of -0.500; two failing families have support radius $\sim d_0$ (Taylor-linearisation breakdown). K-uniform rate holds asymptotically in the natural physical regime (support radius $\gg d_0$). Together with Computation 70 + §18.1.6 Closures (a)–(c), closes the reduced problem at proof-detail level for the i.i.d. surrogate. Stands in for: the emergence-

determined ρ_n, by the same i.i.d. uniform cloud as Computation 70.

Computation 84 · §12.2.2 · [NUMERICAL, superseded by the partition-function route, Computations 87–92/100]

Multi-generator CC inner-fluctuation route for the Z^2 identification.

Extends Computation 74's single-generator inner fluctuation to K coherent quaternion-pair generators $\omega_k = e_{2k-1}e_{2k}$. Computes $D_\phi^2 = D + 2K\phi^2$ and the ϕ^4 coefficient. Matching target gives $y/\Lambda = \sqrt{2K/D}$, requires $K = D/(2\sqrt{2})$ (non-integer) for D-independent target $2^{-1/4}$. Bernoulli-activated extension: $y/\Lambda = \sqrt{p}$ requires $p = 1/\sqrt{2}$, not derivable from $\mu_{\text{site}} = 1/2$. The multi-generator inner-fluctuation route does not deliver the target at a structurally natural value of K or p; the partition-function-level identification of Computations 87–92 takes a different substrate-side route.

Computation 85 · §12.2.2 · [NUMERICAL, superseded by the partition-function route, Computations 87–92/100]

$\mathbf{A}_F$-internal Higgs inner fluctuation for the Z^2 identification.

Places the Higgs in the canonical $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ finite-internal factor (Connes–Chamseddine–Marcolli SM construction) rather than the Clifford volume element. The CCM moment-extraction $\lambda \sim (a/\pi^2) f_0$ reduces to the same identification problem at the moment-extraction level; no new computational test. Identifies the structural obstacle made explicit in Computation 86: the heat-kernel formalism presupposes a continuous spectrum, which the discrete substrate triple lacks. The partition-function-level identification of Computations 87–92 provides the substrate-side route that does not depend on the heat-kernel formalism.

Computation 86 · §12.2.2 · [PROOF, superseded by the partition-function route, Computations 87–92/100]

Seeley–DeWitt heat-kernel analysis on the discrete substrate triple.

Examines the Seeley–DeWitt heat-kernel expansion $\mathrm{Tr}\, f(D^2/\Lambda^2) \sim \sum_n f_n \Lambda^{d-2n} a_n$ for the substrate spectral triple. For the PST substrate at matched scaling: all eigenvalues at $\pm\sqrt{D}$, so $a_0 = 2^D$ (non-zero), $a_n = 0$ for $n \geq 1$ (no continuous-spectrum curvature contributions). The moment structure that delivers λ in continuum CCM is structurally absent. Standard CCM heat-kernel machinery is therefore structurally obstructed for the discrete substrate triple, ruling out the multi-generator CC route flagged by Computation 74. Computations 87–92 take the partition-function-level identification, which does not depend on the heat-kernel formalism.

Computation 87 · §12.3 · [NUMERICAL]

Bernoulli Laplace-transform identification: bypassing the CC inner-fluctuation route.

Bypasses CC inner fluctuation via the asymptotic Bernoulli moment-generating function. $\mathbb{E}_\mu[\exp(c\bar{X})] = ((1 + e^{c/D})/2)^D \to \exp(c/2)$ as $D \to \infty$. Hits e^{-1} at $c = -2$. Numerical convergence verified at $D = 10^4$ with error $\sim 10^{-5}$. Multiple structural readings of $c = -2$ (e.g. $-1/\mu_{\text{site}}$, number of substrate states per bit) compatible; none uniquely forced. Computation 88 fixes the load-bearing reading as the binary-distinction $\mathrm{Cl}(1,0)$ gap $2 = 1 - (-1)$.

Computation 88 · §12.3 · [PROOF]

Binary-gap-tempered Bernoulli MGF identification of e^{-1} at the matched scaling.

The scale $c = -2$ is *structurally forced* by the binary-distinction gap: each substrate exchange τ_i is a self-adjoint involution with eigenvalues $\{+1, -1\}$, so the one-bit Clifford $\mathrm{Cl}(1,0)$ fixes $\beta_{\mathrm{KO}} = 1 - (-1) = 2$ with no freedom (P1's binary alphabet; Computations 100 and 102). Setting $c = -\beta_{\mathrm{KO}} = -2$ gives $e^{-1} = \mathbb{E}_\mu[\exp(-\beta_{\mathrm{KO}}\bar{X})]$ asymptotically. The total KO-dimension $\mathrm{KO}_{\mathrm{total}} = 4 + 6 = 10 \equiv 2 \pmod 8$ shares this value but is structurally independent (an additive residue in $\mathbb{Z}/8\mathbb{Z}$, not a spectral width); the coincidence is retained only as the "binary-gap-tempered" label, never as an input. All three ingredients structurally derived (μ_{site} from P1, the binary-distinction gap 2 from P1, asymptotic limit from Cramér–Bernoulli). Second independent substrate-side derivation of e^{-1}; complements Computation 62's spectral-action route.

Computation 89 · §12.3 · [PROOF]

Z^2 bridge step 1: substrate Higgs Hamiltonian.

Derives the substrate Higgs Hamiltonian $H_{\mathrm{Higgs}}(C)$ structurally from three constraints: (i) additivity (Bernoulli sites independent $\Rightarrow H = \sum_a h_a(B_a)$), (ii) matched scaling (per-site energy $1/D$), (iii) uniformity (P1 symmetry $\Rightarrow h_a$ uniform). Together force $H_{\mathrm{Higgs}}(C) = (1/D)\sum_a B_a = \bar{X}(C)$, no free parameters. The binary-gap-tempered partition function is then $Z_H(\beta_{\mathrm{KO}}) = \mathbb{E}_\mu[\exp(-2\bar{X})] \to e^{-1}$ asymptotically. Substrate side of bridge closed structurally. Numerical verification at $D = 10, 100, 1000$.

Computation 90 · §12.3 · [EMPIRICAL]

Z^2 bridge step 2: three SM-side candidate formulations.

Reduces the SM-side bridge to a choice among three concrete candidate formulations: (I) Wilsonian effective-coupling RG, integrate substrate fluctuations between M_* and Λ; (II) spectral-action variation, naive form gives e^{-2} rather than e^{-1}, fails; (III) substrate-measure invariance, $V_{\text{eff}}(C)/T_{\text{KO}} = \beta_{\text{KO}}\bar{X}(C)$ at matched scaling. Numerical: predicted $\lambda_{\text{SM}}(M_*) = e^{-1}/4 \approx 0.0920$ vs observed 0.0927 (Buttazzo 2013), ratio 1.008. Formulation (II) does not close in naive form; (I) and (III) are taken up as concrete research directions in subsequent work. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 91 · §12.3 · [EMPIRICAL, superseded by Computation 123]

Z^2 empirical-precision test: two-loop SM RGE.

Tests the empirical tightness of $Z^2 \approx e^{-1}$ via SM RGE precision. Simplified one-loop running gives $\lambda_{\text{SM}}(M_*) \approx 0.087$ (5% below target); compact two-loop ≈ 0.085 (7% below). The book's 0.8% near-coincidence claim depends on *full* Buttazzo-level SM RGE precision (two-loop running with explicit threshold corrections at m_t). Sensitivity analysis: variation of $m_h \pm 1$ GeV, $m_t \pm 2$ GeV gives Z^2/e^{-1} in range 0.89–0.97 at simplified one-loop. All readings consistent with structural identity within SM RGE truncation uncertainty. Substrate-side derivation (Computation 89) remains exact and RGE-truncation-independent. *Superseded by Computation 123*, which performs the proper two-loop running with a full input-uncertainty budget, $\lambda_{\text{SM}}(M_*) = 0.0927 \pm 0.0007$ (m_t/m_h-dominated), a rigid 1.0σ match to $\mathrm{e}^{-1}/4 = 0.09197$ that converges monotonically with loop order, confirming the simplified 0.085–0.087 readings here were

truncation artefacts, exactly as anticipated; and in turn by Computation 202, which computes the number at three loops on current inputs.

Computation 92 · §12.3 · [PROOF]

Z^2 bridge formalisation: seven-step structural proof sketch.

Reduces the $Z^2 = e^{-1}$ closure to a single Wilsonian-matching premise via a seven-step structural derivation: (1) PST EFT below M_* = pure SM (eq. (127)); (2) Wilsonian matching at M_*; (3) bare $\lambda_{\text{PST}}^{\text{bare}} = b = 1/4$ (§1.5); (4) substrate UV at $\Lambda = \sqrt{D}$ contributes vacuum fluctuations; (5) fluctuations integrated over μ give $Z_H(\beta) = \mathbb{E}_\mu[\exp(-\beta H_{\text{Higgs}})]$ with $H_{\text{Higgs}} = \bar{X}$ (Computation 89); (6) $\beta_{\text{KO}} = 2$ forced by the binary-distinction one-bit Clifford $\text{Cl}(1,0)$ gap $1-(-1)$ (Computations 88, 100, 112), the KO-mod-8 coincidence retained only as a label, never as an input; (7) $\lambda_{\text{SM}}(M_*) = b \cdot Z_H(\beta_{\text{KO}}) \to e^{-1}/4$ structurally, matching observed 0.0927 at 0.8%. Steps (1)–(4), (6) established in existing PST framework; step (5) by Computation 89; step (7) follows. Elevates Z^2 closure status from "open conjecture" to "structural proof sketch established." The seven steps are subsequently formalised in the partition-function-level Chamseddine–Connes correspondence framework of §12.5; the substrate-side content of Bridge Premise (B) is closed at proof-detail level from P1–P3 by Computation 100 (matched-scaling cutoff identification via tensor-product factorisation forced by P1's Bernoulli product measure) and Computation 98 (mode-shell residual via the discrete Wilson block-spin RG fixed-point theorem); the matching identification of the substrate factor with the SM-side coupling ratio (bridge premise B itself) is subsequently reduced within PST to a field-strength-renormalisation identity (substrate side closed; SM-side matching open) by the wave-function-renormalisation route and the unique-soft-mode theorem (Computations 110–116). *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\text{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on

the 2024 Review).

Computation 93 · §12.2 · [PROOF]
$M_* = 4\pi m_h\sqrt{2/3}$ **off-by-one: fourth-resolution candidate via substrate-vs-emergent distinction.**

Re-examines resolution (b) of Computation 69 under a foundational reading: the substrate-side Boolean zero mode $S = \emptyset$ is the substrate's constant-mode profile, *not* the emergent SM Higgs doublet (captured separately by $C_H = 6\lambda_{\rm PST} = 3/2$ from inner fluctuations on $M \times F$). Under this substrate-vs-emergent distinction, including the zero mode in N_B does *not* double-count the doublet, and total $\delta m_h^2 = (M_*^2/16\pi^2)(C_H + N_B^{\rm full} - N_F) = (3/2)M_*^2/(16\pi^2)$, recovering $m_h = M_*\sqrt{6}/(8\pi)$ as parameter-free one-loop derivation. Identifies the fourth-resolution candidate; full validation via projection-chain analysis follows in Computation 97.

Computation 94 · §12.5 · [NUMERICAL, superseded by Computations 95, 100]
Bridge Premise (B) fluctuation-matching $T_{\rm KO}$ (research-direction sketch).

Attempts to derive Bridge Premise (B) via formulation III with fluctuation-matching prescription $\langle(\bar{X} - 1/2)^2\rangle_\mu = \langle h^2\rangle_{{\rm thermal},T}$, giving the matched-scaling temperature $T_{\rm KO} = \mu^2(M_*)/(4D)$. Combined with $\beta_{\rm KO} = 2$ delivers the testable prediction $\mu^2(M_*) = 2D \cdot M_*^2$. Acknowledged as dimensionally suspect ($\bar{X}$ dimensionless vs h with GeV units; the prediction implies hierarchy violation $\mu \gtrsim M_*$ for $D \gtrsim$ few). Held as research-direction sketch rather than a closure; superseded by Computation 95 reduction and ultimately Computation 100 closure.

Computation 95 · §12.5 · [PROOF]

Bridge Premise (B) reduced to structural claim (***) via substrate-Wilsonian RG.

Reduces Bridge Premise (B) to a single sharper structural claim (***): the substrate Wilsonian RG flow generator equals the Boltzmann partition function $\exp(-\beta_{\mathrm{KO}} \cdot H_{\mathrm{Higgs}}(C))$ on substrate configurations, the partition-function-level analogue of the heat-kernel RG flow generator $\exp(-tD^2)$ for the continuum spectral triple. If (***) is granted, the bridge identification follows structurally: $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z_H(\beta_{\mathrm{KO}}) = e^{-1}/4 \approx 0.0920$. Provides the residual-gap framing subsequently closed by Computations 98 (mode-shell) and 100 (matched-scaling cutoff).

Computation 96 · §12.5 · [NUMERICAL]

Polchinski-discrete derivation of claim (***): partial advance.

Polchinski-style derivation attempted on the discrete substrate. The partition-function form is structurally established via configuration-Hamiltonian Boltzmann averaging $Z_D(\beta_{\mathrm{KO}}) = ((1 + e^{-\beta_{\mathrm{KO}}/D})/2)^D \to e^{-1}$ as $D \to \infty$. Identifies three rigorous-proof items as residual gap: (a) formal exact RG flow equation on discrete substrate triple, (b) uniqueness of Boltzmann form at matched scaling, (c) cutoff-function identification with asymptotic $D \to \infty$ limit. Partial advance; (a) and (b) closed by Computation 98, (c) by Computation 100.

Computation 97 · §12.2 · [PROOF]

M_* off-by-one: projection-chain validation CLOSES the item.

Validates Computation 93's substrate-vs-emergent distinction via explicit PST projection-chain analysis $\psi \to \Pi(\psi) \to H_{\mathrm{doublet}}$. The substrate Walsh zero

mode $\chi_\emptyset$ maps under Π to the constant function on $M = \mathbb{R} \times S^3$ (vacuum direction), structurally distinct from the emergent SM Higgs doublet (4 components from inner fluctuations on $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$). Including the zero mode in N_B does *not* double-count $C_H = 6\lambda = 3/2$. Total $\delta m_h^2 = (3/2)M_*^2/(16\pi^2)$, restoring $m_h = M_*\sqrt{6}/(8\pi)$ as parameter-free one-loop derivation. The off-by-one of Computations 67 and 69 is closed at proof-detail level; the "one-loop-complete / parameter-free" reading of $M_* = 4\pi m_h\sqrt{2/3}$ is restored.

Computation 98 · §12.5 · [PROOF]

Bridge Premise (B) items (a)+(b) CLOSED via mode-shell block-spin RG.

Closes items (a), (b) of Computation 96 at proof-detail level via explicit discrete Wilson block-spin RG. Key observation: $H_{\text{Higgs}} = \bar{X}$ is supported on the lowest two Walsh-mode shells $V_0 \oplus V_1$ only, with explicit expansion $\bar{X} = (1/2)\chi_\emptyset - (1/(2D))\sum_{|S|=1}\chi_S$. Because $\bar{X}$ has no $|S| \geq 2$ content, the block-spin RG step from level $m+1$ to m factorises trivially: $\bar{X}(\psi_{\leq m}+\psi_{m+1}) = \bar{X}(\psi_{\leq m})$ for every $m \geq 1$, so the Boltzmann action $\beta_{\text{KO}}\bar{X}$ is a TRIVIAL FIXED POINT (Theorem 1). Uniqueness follows from Computation 89 boundary condition. Residual item (c) remains as the matched-scaling identification (subsequently closed by Computation 100).

Computation 99 · §12.5 · [NUMERICAL, superseded by Computation 100]

Item (c) via thermal-QFT framing (EXPLORATORY, superseded by Computation 100).

Exploratory closure of item (c) via thermal-field-theory framing: postulates $\beta_{\text{KO}} = \text{KO mod } 8$ as substrate's intrinsic inverse temperature in the partition-function-level CC framework. Standard thermal-QFT IR coupling formula $\lambda_{\text{eff}} = \lambda_{\text{bare}} \cdot Z(\beta)$ delivers $\lambda_{\text{SM}}(M_*) = b \cdot Z_H(\beta_{\text{KO}}) = b \cdot e^{-1}$. The "KO-thermal

principle" is at the same level as the spectral-action principle in standard CC, a framework postulate beyond P1–P3. Held as exploratory, not as a P1–P3-only closure: PST's claim is that all structure derives from P1–P3 alone, so a CC-framework postulate weakens that claim. SUPERSEDED by Computation 100, which closes (c) from P1–P3 alone via tensor-product factorisation, eliminating the thermal-postulate requirement.

Computation 100 · §12.5 · [PROOF]

Substrate-side spectral-action cutoff identification closed from P1–P3 via tensor product.

Closes the substrate-side content of item (c) of Computation 96 from P1–P3. P1's Bernoulli independence forces tensor product $L^2(\mathcal{P}(D), \mu) = \bigotimes_i L^2(\{0,1\}, \mu_i)$. The Boolean Laplacian $\Delta = \sum_i (1 - \tau_i)$ with one-bit Clifford $\mathrm{Cl}(1,0)$ gives each $(1 - \tau_i)$ eigenvalues $\{0, 2\}$. For $\mathrm{Tr}\, f(\Delta/\Lambda^2)$ to be CONSISTENT with the tensor product structure (i.e., to factor as a product of one-bit traces over independent bits) f must satisfy $f(x + y) = f(x)f(y)$, unique smooth solution $f(x) = \exp(-x)$. Forced by P1's product structure, not postulated. At matched scaling $\Lambda^2 = D$: $(1/2^D)\,\mathrm{Tr}\, e^{-\Delta/D} = ((1 + e^{-2/D})/2)^D \to e^{-1}$ by binomial product convergence. The factor "2" in $\exp(-2/D)$ comes from one-bit Clifford structure, *not* KO mod 8 thermal interpretation (numerically equal but structurally independent). Computation 99's KO-thermal postulate SUPERSEDED. *The matching identification of this substrate-side factor with the SM-side coupling ratio* $\lambda_{\mathrm{SM}}(M_*)/b$ (bridge premise (B) itself) *is subsequently reduced within PST to a field-strength-renormalisation identity (substrate side closed; SM-side matching open)* by the wave-function-renormalisation route (Computations 110–116); $\lambda_{\mathrm{SM}}(M_*) = b \cdot e^{-1} = 0.0920$ matches observed 0.0927 at 0.8% Buttazzo-RGE precision (5–7% at one-loop, Computation 91). *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008

on the 2026 Review, 1.4σ from $e^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 101 · §12.5 · [PROOF]

Bridge premise (B) sharpened: Z^2 as wave-function renormalisation Z_ϕ^2.

Addresses the structural concern that standard Wilsonian threshold matching of a quartic is additive in loop corrections, not multiplicative. Reinterprets the multiplicative matching $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z_H(\beta_{\mathrm{KO}})$ as a wave-function renormalisation of the substrate modal field ψ to the SM Higgs ϕ at matched scaling, with the substrate-side partition function $Z_H(\beta_{\mathrm{KO}}) \to e^{-1}$ playing the role of Z_ϕ^2 via standard QFT decomposition $Z^2 = Z_{\Gamma 4}/Z_\phi^2$. Reduces (B) to two well-posed sub-questions: (1) is the matched-scaling Π's path-integral Jacobian equal to $\exp(-\beta_{\mathrm{KO}} H_{\mathrm{Higgs}}(C))$? (2) why is the vertex renormalisation trivial ($Z_{\Gamma 4} = 1$)? Both are tractable extensions of Computation 73's matched-scaling machinery; closing both delivers (B) automatically via standard renormalisation. Status: structural reduction of (B)'s open content from a non-standard hypothesis to two standard renormalisation questions. Both sub-questions are subsequently closed: the multiplicative (not additive) reading is the correct one (the additive-Jacobian route is excluded, Computation 103) and the vertex renormalisation $Z_{\Gamma 4} = 1$ holds to all orders (Computation 119); (B) is thereby reduced (Computations 110–119).

Computation 102 · §12.5 · [PROOF]

Item 1.2 resolution: the two "2"s are structurally independent.

Addresses the concern that the identification $c = -(\mathrm{KO}_{\mathrm{total}} \bmod 8) = -2$ in the Bernoulli MGF is asserted, not motivated. Computation 100 already derived the "2" from one-bit Clifford structure $\mathrm{Cl}(1,0)$ (eigenvalue range of

$(1 - \tau_i)$), independent of KO mod 8. This computation makes the structural-independence explicit: $\mathrm{KO}_{\mathrm{total}} \bmod 8 = 2$ (from spacetime + internal split $4 + 6 = 10$) and one-bit Clifford eigenvalue range = 2 (from $\tau_i^2 = 1$) trace to different parts of the substrate construction and are numerically coincident but structurally unrelated. The honest framing is that Computation 100's derivation uses ONLY one-bit Clifford structure; KO mod 8 = 2 is a parallel substrate fact, not an input to the matched-scaling exponent. Resolves item 1.2 as a derivation question; the remaining open content (whether partition-function-level CC has standing in the spectral-action literature) is an external-validation / literature-search task, not new research.

Computation 103 · §12.5 · [NUMERICAL, superseded by Computations 110–116]

Bridge premise (B) attack: explicit path-integral Jacobian, honest negative result.

Attempts Computation 101's sub-question (1): is the matched-scaling Π's path-integral Jacobian equal to $\exp(-\beta_{\mathrm{KO}} H_{\mathrm{Higgs}}(C))$? Two natural interpretations both fail: (A) the linear-map determinant of $\Pi : V_0 \oplus V_1 \to \mathrm{SM}_{\mathrm{low}}$ is configuration-independent (constant); (B) the Cramér–Bernoulli LDP near the vacuum $\bar{X} = 1/2$ gives a Gaussian weight in $(\bar{X} - 1/2)$, not exponential in $\bar{X}$ (numerically verified $I(x) = x \ln(2x) + (1 - x) \ln(2(1 - x)) \approx 2(x - 1/2)^2$ near 1/2). The multiplicative matching $\lambda_{\mathrm{SM}}(M_*) = b \cdot Z_H(\beta_{\mathrm{KO}})$ is therefore *not* derivable from a path-integral change-of-variable. Confirms that the *Jacobian* route to (B) does not close. Sub-question (2) [$Z_{\Gamma 4} = 1$] is satisfied at tree level (standard CC inner fluctuations). Net: this is a negative result for the measure-Jacobian reading specifically; it does *not* obstruct (B), because (B) is a *wave-function* renormalisation (multiplicative), not a coupling-matching Jacobian (additive). The Gaussian found here is precisely what the product-structure argument later excludes for the field-strength kernel (Computation 111), and (B) is subsequently reduced within PST to a field-

strength-renormalisation identity (substrate side closed; SM-side matching open) by the wave-function route (Computations 110–116).

Computation 104 · §12.5 · [PROOF, partially retracted, superseded by Computations 110–119]

Bridge premise (B) factored: substrate side derived, SM side inherited from standard CC literature.

Attempts to extend Mosco convergence from Dirichlet forms (which §18.1 establishes) to spectral actions, with the goal of deriving bridge premise (B) at the same level Mosco delivers the SM kinetic term. The structural analysis shows (B) FACTORS into two parts: (B-substrate) $(1/2^{|D|})\,\mathrm{Tr}\, f(\Delta/|D|) \to e^{-1}$, which is CLOSED at proof-detail level from P1–P3 by Computation 100; and (B-SM) the SM-side spectral-action identification of the matched-scaling coefficient with $\lambda_{\mathrm{SM}}(M_*)$, which *is* the standard Chamseddine–Connes spectral-action principle [49, 50] applied to PST's substrate spectral triple. This is an ESTABLISHED FRAMEWORK with 30+ years of literature precedent, *not* a PST-specific postulate. PST extends standard CC by providing a pre-causal foundation BENEATH the spectral triple (P1–P3 deliver the substrate spectral triple itself, Computations 3, 19); it does not invent the SM-side framework. Whether the partition-function-level CC framework has standing in the spectral-action literature: yes, the framework PST uses *is* standard CC, with the partition-function-level adaptation being PST's specific way of handling the discrete substrate where heat-kernel collapses (Computations 85–86). PST's foundational claim refines to "P1–P3 + standard CC spectral-action framework": the substrate-side contribution is the genuinely new PST result (Computation 100, derived from P1–P3), and the SM-side framework is inherited from established CC literature.

This inheritance framing is subsequently superseded: the matching (B) is reduced *internally* to PST, in the substrate limit, to a field-strength-renormalisation identity whose internal content is fully derived (field-strength

and vertex, Computations 110–119), without inheriting an SM-side CC postulate; the foundational claim is then "P1–P3 + matched scaling", with the bridge resting on the emergence premise F4 (§12.6.2.1) + the RGE test rather than an inherited framework.

Computation 105 · §12.5 · [NUMERICAL]

Two-layer one-loop diagram (no double counting of M_* zero modes).

Addresses the concern of potential double-counting between the substrate $\chi_\emptyset$ singleton-cylinder order parameter and the emergent SM Higgs field C_H from Mosco convergence. The PST partition function factorises as $Z_{\mathrm{PST}}(M_*) = Z_{\mathrm{sub}}(M_*, \Lambda{=}\sqrt{D}) \cdot Z_{\mathrm{EFT}}(M_*, \mathrm{IR})$, with the two layers integrating over disjoint momentum ranges $[M_*^2, D]$ and $[m_h^2, M_*^2]$ that meet only at the single matching scale M_* (measure zero in both integrals). Three independent non-double-counting criteria are established: (i) disjoint momentum support; (ii) distinct dynamical variables ($\chi_\emptyset$ on $\mathcal{P}(D)$ vs C_H on $M = \mathbb{R} \times S^3$); (iii) distinct counterterm structures (the matching is a boundary condition, not a loop integration). Numerical verification: the substrate full-trace partition function (Computation 100) decomposes exactly as $Z^{\mathrm{full}}_{\mathrm{sub}}(D) = Z^{\mathrm{IR}}_{\mathrm{sub}}(D, k_*) + Z^{\mathrm{UV}}_{\mathrm{sub}}(D, k_*)$ across the M_*-shell partition with $|Z^{\mathrm{full}} - (Z^{\mathrm{IR}} + Z^{\mathrm{UV}})| < 10^{-14}$ for $D \in \{10, 20, 50, 100\}$. Closes the double-counting question at the diagram level.

Computation 106 · §14.2 · [NUMERICAL, supersedes the earlier PROOF]

Colour does not preserve the generation sectors, and the earlier block-diagonality claim is withdrawn.

An earlier version claimed every $g \in \mathrm{SU}(3)_c \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ preserves each generation sector G_k, and concluded that tree-level flavour-changing neutral currents are absent as a structural theorem. *That is withdrawn.* Its part (1) asserted that $\mathrm{SU}(3)_c$ acts within a single Fano sub-algebra through $\hat{\tau}$; but

$\mathrm{Stab}_{G_2}(\hat{\tau})$ *fixes* $\hat{\tau}$, and each G_k's octonionic factor is one complex line, which cannot carry a colour triplet. Its numerical check verified a Kronecker identity: $\not\vdash_{\mathrm{gen}} \otimes X$ is block-diagonal in the generation index for *every* X, so the check could not fail. *What this computation now does*: builds $\mathrm{Der}(\mathbb{O})$ (dimension 14), takes the subalgebra annihilating $\hat{\tau}$ (dimension 8, the $\mathfrak{su}(3)$), exponentiates to genuine octonion automorphisms, and *measures* how much each generation sector leaks. The leakage is $O(1)$, and the subalgebra preserving all three Fano lines through $\hat{\tau}$ is two-dimensional, the maximal torus. *Conclusion*: parts (2) and (3) stand, $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ being generation-blind by the tensor-product structure; colour is not. Tree-level FCNC absence is not established, and would additionally require a statement about the *mass* eigenbasis that §14.3 places in the contingent layer.

Computation 107 · §12.2 · [NUMERICAL]

Full-SM extension of the M_* boson-fermion cancellation: sector-blindness sub-result.

The substrate-level binomial cancellation $N_B = N_F = 2^{D-1}$ is SECTOR-BLIND (invariant under any partition of the 2^D Walsh modes; numerically verified across $D \in \{4, \ldots, 20\}$; superseded by Computation 120, which settles the extension in the negative). This sector-blindness does *not* transfer the cancellation to the SM top/W/Z sectors: by the disjoint-layer structure of Computation 105 the substrate cancellation lives in the UV shell $[M_*, \sqrt{D}]$ while the physical top/W/Z loops run in the emergent EFT shell below M_*, so the substrate identity cannot cancel the EFT-layer loops. The "coupling universality at M_*" premise ($c_S = 1$ for every Walsh mode) is contradicted by the observed Yukawa hierarchy, which the book assigns to $T(C)$-contingency (§14.3). The empirical SM Veltman residual at M_* is ≈ -3 (top-dominated, negative), not the $+m_h^2$ a clean cancellation would leave. The sector-blindness of the substrate count is a fact; it does not close the full-SM extension, which does not exist within PST (Computation 120), so M_* is irreducibly a

scalar-sector prediction.

Computation 108 · §12.2 · [PROOF]

Inner-fluctuation sign assignment for SM gauge bosons is BOSONIC (sub-result; does not close the full-SM extension).

Verifies that the substrate $|S|$-parity boson/fermion sign assignment (CAR realisation, §11.2) carries through to A_F inner-fluctuation gauge bosons that are not themselves Walsh modes. Three independent readings agree: (R1) algebraic, Connes 1-forms $A = \sum_i a_i[D, b_i]$ for $a_i, b_i \in A_F$ are operators of EVEN Clifford grade; (R2) substrate-side lift, A_F is generated by $\mathrm{Cl}_{\mathrm{even}}(\mathcal{P}(D))$ under the CAR realisation, so every inner fluctuation lifts to an $|S|$-even substrate configuration; (R3) one-loop diagrammatic, the standard QFT vector-boson loop sign (+) is inherited unchanged by the CC inner-fluctuation construction. Dof count: bosonic inner-fluctuation dofs on A_F (off-shell) are $\mathrm{U}(1)_Y$: 4, $\mathrm{SU}(2)_L$: 12, $\mathrm{SU}(3)_c$: 32, Higgs doublet: 4, total 52; the substrate $|S|$-even Walsh subspace has dimension 2^{D-1} accommodating these at $D \geq 6$. *Status:* the bosonic sign assignment is a correct sub-result, but it does *not* close the full-SM-extension question. The full-SM cancellation fails for the layer-disjointness reason of Computation 107 (the substrate cancellation does not reach the EFT-layer top/W/Z loops), independent of the inner-fluctuation sign. That question was subsequently settled in the negative by Computation 120 (the full-SM extension does not exist within PST); this sign sub-result stands.

Computation 109 · §12.2 · [PROOF]

Higgs-quartic normalisation + O(4) one-loop self-energy: two technical verifications.

Verifies two technical findings on the Higgs-sector derivation at proof-detail level. *(A) Normalisation:* the book commits to the *doublet convention*, $\lambda_{\mathrm{PST}} =$

$\frac{1}{4}$ is the coefficient of $(\Phi^\dagger\Phi)^2$, the same convention as the SM $V_{\rm SM} = \lambda_{\rm SM}(\Phi^\dagger\Phi)^2$, with $\langle\psi,\psi\rangle \leftrightarrow \Phi^\dagger\Phi$. Under the radial reduction $\langle\psi,\psi\rangle \to \psi^2$ the postulate $\mathcal{F} = \int[-\varepsilon\langle\psi,\psi\rangle + (1/4)\langle\psi,\psi\rangle^2 + c|\nabla\psi|^2]\mathrm{d}\mu$ collapses to the scalar form $-\varepsilon\psi^2 + (1/4)\psi^4$ (eq. (24)). The field-normalisation ratio $Z^2 = \lambda_{\rm SM}(M_*)/\lambda_{\rm PST}$ is then a ratio of two same-convention couplings. *(B) Vertex contraction:* explicit O(N) vertex contraction $V_{ijkl} = 2\lambda(\delta_{ij}\delta_{kl} + \delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk})$ closed into the one-loop seagull $(1/2)V_{ijkl}\delta_{kl} = (N+2)\lambda\,\delta_{ij}$ delivers the standard O(N) result $(N+2)\lambda$ rather than $N\lambda$. For $N = 4$ the Higgs self-loop coefficient is $C_H = 6\lambda_{\rm PST} = 3/2$, giving $M_* = 4\pi m_h\sqrt{2/3} \approx 1285$ GeV. Broken-phase cross-check: 3λ (physical Higgs loop) $+3\lambda$ (three Goldstones) $= 6\lambda$ confirms the symmetric-phase result. This computation propagates corrections through eq. (128), the M_* value, the naturalness coefficient, and the affected Computations 67, 69, 93, 97, 107, 108.

Computation 110 · §12.5 · [PROOF]

Bridge premise (B): wave-function-renormalisation route.

Recasts (B), following Computation 101, as a field-strength (wave-function) renormalisation Z_ϕ^2, which acts multiplicatively on the quartic ($\lambda \to \lambda Z_\phi^2$), the form (B) requires, sidestepping the additive-matching obstruction Computation 103 found for the coupling-Jacobian reading. The route first sets up the discrete Wilson–Polchinski flow on the Boolean lattice with the P1-forced cutoff $f = \exp$ and matched scaling $\Lambda^2 = D$, reducing to the low Walsh shells $V_0 \oplus V_1$ where $\bar{X}$ lives (Computation 98). It then verifies that the field-strength factor is $Z_\phi^2 = Z_H(\beta_{\rm KO}) \to e^{-1}$, so $\lambda_{\rm SM} = b\,Z_\phi^2 \to e^{-1}/4$. Sets up the route; the crux (the form and value of the kernel) is closed in Computations 111–116.

Computation 111 · §12.5 · [PROOF]

Bridge premise (B): the field-strength kernel is forced into $\exp(-\beta\bar{X})$ form.

The field-strength dressing must respect P1's product measure (the matched-scaling cutoff factorises over bits, Computation 100), so the per-configuration kernel $K(C) = \prod_a g(C_a)$ is a product over bits. A product kernel over Boolean bits is necessarily exp of a quantity *linear* in $\bar{X}$, i.e. $\propto \exp(-\beta\bar{X})$, the same form the substrate side evaluates to e^{-1}. The Gaussian $\exp(-\alpha(\bar{X} - 1/2)^2)$ that obstructs the Jacobian route (Computation 103) is thereby excluded: it correlates the bits, breaking the product structure (verified: bit–bit covariance 0 under $\exp(-2\bar{X})$, nonzero under the Gaussian). The wave-function route is therefore unobstructed, in contrast to the Jacobian route.

Computation 112 · §12.5 · [PROOF]

Bridge premise (B): $\beta_{\mathrm{KO}} = 2$ is the binary-distinction gap.

The load-bearing exponent is the spectral range of a single binary distinction, $2 = 1 - (-1)$: each exchange τ_i is a self-adjoint involution with spectrum $\{+1, -1\}$, so $(1 - \tau_i)$ has nonzero eigenvalue 2, forced by P1's binary alphabet with no freedom (equivalently the one-bit Clifford $\mathrm{Cl}(1, 0)$ range). q-ary comparison: only $q = 2$ gives the self-adjoint involution P1 requires, and e^{-1} is the signature of the binary gap ($Z_H \to e^{-\beta/2}$ lands on e^{-1} only at $\beta = 2$). Dissolves the "two 2's" question: $\mathrm{KO}_{\mathrm{total}} \bmod 8 = 2$ shares the value but is a gauge/spacetime-sector fact, never an input.

Computation 113 · §12.5 · [PROOF]

Bridge premise (B): the $Z_\phi \leftrightarrow$ trace identification, total-reduction residual.

Tests the field-strength↔matched-scaling-trace identification ab initio. The order-parameter two-point residue (single $|S| = 1$ shell) gives a single-mode factor, not the full e^{-1}, unless all substrate modes are folded in. The matched-scaling trace and the tempered order-parameter expectation $\mathbb{E}_\mu[\exp(-2\bar{X})]$ are verified exactly equal, so the identification holds iff the substrate→Higgs reduction is *total*. Isolates the residual as one concrete property (total reduction), closed by Computations 114–116.

Computation 114 · §12.5 · [PROOF]

Bridge premise (B): the standard-EFT normalisation is excluded; embedding-norm reading.

Two field-theoretic normalisations of the reduction: (A) the total-reduction embedding norm of the symmetric Higgs vacuum $= \mathbb{E}_\mu[\exp(-\beta_{\mathrm{KO}}\bar{X})] \to e^{-1}$, verified exactly equal to the tempered vacuum amplitude; (B) the standard-EFT two-point residue $\mathrm{Var}_\mu(\bar{X}) = 1/(4D) \to 0$. The bridge needs (A); route (B), the Higgs as a light field whose two-point function is corrected by heavy loops, gives $\lambda_{\mathrm{SM}} \to 0$ and is excluded by the finiteness of the observed quartic. The Higgs is an emergent projected field, normalised by its embedding norm, not a fundamental light field.

Computation 115 · §12.5 · [PROOF]

Bridge premise (B): unique-soft-mode theorem, total reduction.

Past the LG threshold (P3) the order parameter $\bar{X}$ is the unique soft mode. Lemma: $\bar{X} = (1/D)\sum_a C_a$ is linear in the bits, so the LG Hessian is rank one (the all-ones matrix), softening exactly the S_D-singlet direction ($= \bar{X}$) and leaving the $D-1$ standard-rep modes and all $|S| \geq 2$ shells at the binary gap. Verified: rank $= 1$, one soft eigenvalue, $D-1$ pinned at the gap; the threshold sweep softens $\bar{X}$ alone. Establishes total reduction (Theorem F10, Gaussian order), so the bridge factor is the full embedding norm e^{-1}.

Computation 116 · §12.5 · [PROOF]

Bridge premise (B): all-orders decoupling theorem.

Upgrades the unique-soft-mode result to all orders. Because the LG potential is a function of the linear collective coordinate $\bar{X}$ alone, every derivative tensor of V is the totally symmetric all-ones tensor $V^{(n)}/D^n$, with zero matrix element on the non-collective subspace at every order: the standard-rep and $|S| \geq 2$ modes are free fields, decoupled, pinned at the binary gap, with no V-vertex from which any correction could be built. Verified exactly with the full $\exp(-V)$: the V-induced bit correlation is $O(1/D)$, the standard channel sits at the free binary-gap value with $O(1/D)$ deviation, and only the singlet $\bar{X}$ responds to V. The spectral gap is protected non-perturbatively in the substrate limit; the residual is the standard $O(1/D)$ finite-size correction.

Computation 117 · §12.5 · [PROOF]

Bridge premise (B): the embedding-norm selection is structural, not empirical.

Removes the empirical crutch in the field-strength selection. The two-point-residue reading sets $Z_\phi^2 = \mathrm{Var}_\mu(\bar{X}) = 1/(4D) \to 0$; that vanishing is the *law of large numbers*, $\bar{X}$ is an intensive average of the i.i.d. bits (P1), so it self-averages to a sharp value, and a quantity the LLN sends to zero cannot be a field-strength (it gives a free emergent Higgs, $\lambda_{\mathrm{SM}} \to 0$). The embedding norm is the $O(1)$ object $\to e^{-1}$, the only candidate that normalises a canonical emergent field. Since the emergent EFT below M_* is the *interacting* Standard Model (F4), an interacting Higgs is required, hence $Z_\phi^2 = O(1)$, hence the embedding-norm reading, forced by P1's LLN and F4, not by the measured value $\lambda_{\mathrm{SM}} = 0.0927$, which is demoted from an input to the selection to a test of the prediction $b\,e^{-1}$. Non-circular: F4 selects the reading, the substrate gives the value e^{-1}. Physically, below the LG threshold $\langle\bar{X}\rangle$ is a sharp classical condensate and the physical Higgs is the radial excitation, normalised by the embedding norm, not the order parameter's vanishing susceptibility. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 118 · §12.5 · [PROOF]

Bridge premise (B): F10 decoupling as a formal all-orders S_D-equivariant Morse–Bott lemma.

States and proves the F10 decoupling as a formal lemma. Because $F_{\mathrm{int}} = \Phi(\bar{X})$ factors through the *linear* collective coordinate, the chain rule gives $\partial^n F_{\mathrm{int}} = \Phi^{(n)}(\bar{X})/D^n$ times the all-ones tensor at every order. Under S_D the

fluctuation space splits into the singlet line (carrying $\bar{X}$) and the standard representation $W = \{v : \sum_a v_a = 0\}$; an all-ones tensor contracted with any $v \in W$ gives $(\sum_a v_a)(\cdots) = 0$, so the interaction has zero matrix element on W at every order and the W-Hessian is the bare Boolean Laplacian, bounded below by the binary gap ≥ 2. Schur's lemma (S_D-equivariance) forbids singlet–W mixing and any invariant deformation of the W-block. This is the all-orders S_D-equivariant Morse–Bott structure: a soft critical direction (the order parameter) and a uniformly gapped normal bundle. Verified: derivative tensors index-independent at $n = 2, 3, 4$; annihilation of W to machine precision; the singlet softens $4 \to 0$ under the LG curvature while W stays pinned at 2; a random S_D-invariant Hessian has W-eigenvalue independent of the singlet coupling. Residual: the $O(1/D)$ finite-size sum-constraint correction (Computation 116), vanishing as $D \to \infty$. SUBSUMED by Computation 121: the decoupling is closed exactly and non-perturbatively by sufficiency of the empirical mean (this Morse–Bott lemma is its de Moivre–Laplace shadow), with the $O(1/D)$ isolated to the value; Computation 122 shows F8's linearity controls only that $O(1/D)$.

Computation 119 · §12.5 · [PROOF]

Bridge premise (B): vertex renormalisation $Z_{\Gamma 4} = 1$ to all orders, completing the matching's internal content.

The matching ratio $Z^2 = \lambda_{\mathrm{SM}}/b$ decomposes into a field-normalisation piece (the embedding norm $\to e^{-1}$, Computation 117) and a vertex piece $Z_{\Gamma 4}$ (the four-point 1PI correction). Computation 101 fixed $Z_{\Gamma 4} = 1$ only at tree level; this computation closes it to all orders. The matching integrates out the substrate UV modes (the non-collective S_D standard representation W and the gapped $|S| \geq 2$ shells); by the Morse–Bott decoupling lemma (Computation 118), $V = \Phi(\bar{X})$ has zero matrix element on W at every order, and the $n = 4$ case is the Higgs quartic vertex, so the integrated-out modes do not couple to it and generate no vertex correction. Verified: the quartic all-

ones tensor contracted on W legs is machine-zero. Hence $Z_{\Gamma 4} = 1$ to all orders and $Z^2 = e^{-1}$ with no residual internal renormalisation. The only quartic loop corrections are the collective Higgs self-loops, the SM RGE running below M_*, which confirms $b\,e^{-1} = 0.0920$ against the two-loop-run 0.0927 at 0.8% (quantified as a rigid 1.0σ match by Computation 123), not a $Z_{\Gamma 4}$ correction. Net: the matching's internal renormalisation content is complete; the open content of (B) reduces to the emergence premise F4 plus the RGE test, not an internal step. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 120 · §12.2 · [EMPIRICAL]

Full-SM extension of the M_* cancellation settled in the negative.

Answers the open full-SM-extension question (does an EFT-layer top/W/Z cancellation exist, and what partner spectrum would it need?): no, not within PST. The SM Veltman bracket $[6\lambda + (9/4)g^2 + (3/4)g'^2 - 6y_t^2]$ is -3.5 at the EW scale and -3.3 at M_*, top-dominated, with no zero crossing, M_* is not a Veltman-natural full-theory scale. Cancelling the dominant $-6y_t^2$ requires bosonic top-partners (stop-like) at top strength near $M_* \approx 1.3$ TeV; such a sector is collider-disfavoured (LHC excludes scalar top-partners / vector-like quarks to $\sim$ 1.2–1.4 TeV, i.e. at M_*), forbidden by the "no new light states" premise (§15.11), absent from PST's spectrum (F4: emergent EFT = SM, no superpartners), and unreachable by the substrate mechanism (Computation 105: layer-disjoint). Conclusion: $M_* = 4\pi m_h\sqrt{2/3}$ is irreducibly a scalar-sector prediction, exactly as the book scopes it. This is a negative settlement, a definitive limitation on the scope of M_*, not a positive closure; it removes the item from the open register with M_* correctly and finally scoped.

Computation 121 · §12.5 · [PROOF]

Bridge premise (B): the all-orders decoupling closed exactly by sufficiency of the empirical mean (subsumes Computation 118).

Closes the one residual rigour task carried by F10/Computation 118 (the all-orders S_D-equivariant decoupling) by an exact, non-perturbative identity. Theorem F8 makes $H_{\text{Higgs}} = \bar{X}$ exactly the empirical mean of i.i.d. per-site bits, so the interaction $\Phi(\bar{X})$ is a function of a *sufficient statistic*. By Fisher–Neyman factorisation the configurations at fixed $\bar{X} = k/D$ are counted by the interaction-independent multiplicity $\binom{D}{k}$, and the substrate partition function factorises exactly at every finite D: $Z[\Phi] = \sum_k \binom{D}{k} 2^{-D} e^{-\Phi(k/D)}$. The transverse (non-collective) modes enter only through the Φ-independent multiplicity, so $Z_{\Gamma 4} = 1$ (and every n-point Φ-vertex correction) vanishes exactly, with no "higher-order" gap (the factorisation holds at all orders simultaneously). Verified: the factorisation to machine precision for linear/quartic/mixed Φ; the quartic-probe vertex test ($Z_{\Gamma 4} = 1$ to 3×10^{-16}) with a transverse-coupled control that breaks sufficiency (gap 0.30, the test has teeth); the $O(1/D)$ isolated to the value, the decoupling carrying no D-dependence; and the de Moivre–Laplace reduction showing Computation 118's Morse–Bott lemma is the Gaussian shadow. Net: the substrate side of (B) carries only the standard $O(1/D)$ caveat on the value e^{-1}; the single open content of (B) remains the SM-side matching (F4 + RGE; quantified as a rigid, D-independent 1.0σ test by Computation 123). *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 122 · §12.5 · [PROOF]

Bridge premise (B): what the substrate side needs of F8 linearity (decoupling = symmetry; value = vacuum normalisation).

Isolates the role of Theorem F8's linearity ($H_{\text{Higgs}} = \bar{X}$) and finds it bears far less weight than "the whole substrate side hangs on exact linearity". (1) Uniformity alone (the S_D-symmetry that *is* P1) forces $H_{\text{Higgs}} = g(\bar{X})$ for some g, because any symmetric function of binary bits depends only on the count; this is exact and is all that the all-orders decoupling (Computations 118, 121) needs. (2) Linearity enters only via additivity: a single-bit function is necessarily affine, so an additive Hamiltonian is exactly affine in $\bar{X}$ with no $O(1/D)$ remainder; the only route to a nonlinear g is a cross-bit term, an interaction between P1-independent distinctions, additivity is P1 independence read at the level of the energy. (3) The leading value e^{-1} is robust to relaxing linearity: by the LLN $\bar{X} \to 1/2$, so $Z_H(\beta) \to e^{-\beta g(1/2)}$, depending on g only through $g(1/2)$; with the binary gap $\beta = 2$, e^{-1} follows from the single normalisation $g(1/2) = 1/2$. Verified: four g's (linear + nonlinear deformations) all with $g(1/2) = 1/2$ give e^{-1} to leading order, the difference $O(1/D)$; a slope sweep confirms value $= e^{-2g(1/2)}$ exactly in the limit. Net: the substrate-side closure of (B) does not hinge on H_{Higgs} being exactly linear, its structure is exact from symmetry, its value rests on a vacuum normalisation plus the binary-gap temperature, and linearity refines only the $O(1/D)$.

Computation 123 · §12.5 · [EMPIRICAL]

Bridge premise (B): two-loop SM-side error budget and the finite-D rigidity of the e^{-1} matching.

Strengthens the SM-side matching (the single open content of (B)) from a "0.8% agreement" to a quantified, rigid, falsifiable test. (A) SM side: running the SM quartic up to $M_* = 1285$ GeV at two loops gives $\lambda_{\text{SM}}(M_*) =$

0.0927 ± 0.0007, the uncertainty dominated by the top and Higgs masses (parametric ± 0.00068: $m_t \pm 0.00054$, $m_h \pm 0.00040$, $\alpha_s \pm 0.00010$; truncation ± 0.0002), *not* by logarithmic running. The value converges monotonically with loop order ($0.09145 \to 0.09234 \to 0.09270$, steps shrinking), so the match is a genuine two-loop statement, not a low-order accident; the integration is validated by the vacuum instability scale. Against the substrate prediction $\mathrm{e}^{-1}/4 = 0.09197$ the gap is 1.0σ, consistent. (B) Substrate side: the prediction is *rigid*. By eq. (134) the field-normalisation factor is e^{-1} exactly and D-independently (the matched-scaling eigenvalue collapse $\pm\sqrt{D}/\Lambda \to \pm 1$); in the partition-function route the $O(1/D)$ correction would need an implausible $D \approx 63$ to close the gap, against a substrate cardinality $|D| = V/a_n^4$ (site spacing $a_n \to 0$) larger by dozens of orders of magnitude, so there is no finite-D freedom to absorb the residual. Superseded on the SM side by Computation 202: the anchor 0.0927 is not reproduced by its source's own formulas (they give 0.0924 on the 2013 inputs) and the budget counted m_t at ± 0.5 GeV where 2013 had ± 0.76; the current statement is 0.0930 ± 0.0008, 1.4σ on the 2026 Review. Net, as this computation stood: $\lambda_{\mathrm{SM}}(M_*) = \mathrm{e}^{-1}/4$ is a rigid, parameter-free prediction, consistent with the SM at 1.0σ, with the residual living entirely in the m_t/m_h inputs and sharply falsifiable as those (and the SM matching order) tighten.

Computation 124 · §7.8 · [PROOF]

Displacement covariance is positive-semidefinite; the Lorentzian signature is the τ_{null} sign cut, not isotropy.

Backs the clause at (56). The empirical displacement covariance $C^{\mu\nu} = |D|^{-1} \sum_a (\rho(a)^\mu - x^\mu)(\rho(a)^\nu - x^\nu)$ is a sum of outer products, hence positive-semidefinite (all four eigenvalues $\to d_0^2/4 > 0$), whereas the Lorentzian $g = \mathrm{diag}(+1, -1, -1, -1)$ is indefinite, so C cannot equal $(d_0^2/4)g$. Under $O(4)$ isotropy $C \to (d_0^2/4)\delta$ (eigenvalue ratio $\to 1$, off-diagonal $\to 0$ as $N \to \infty$): isotropy fixes only the positive Euclidean magnitude and its scale. The

signature is the τ_{null} sign cut of (30) (Computation 2): across a quantile sweep the timelike fraction tracks the cut exactly while C stays PSD throughout. Confirms (56) is read in the post-signature metric.

Computation 125 · §18.1.6 · [SURROGATE]

Locality of the emergence map $\rho_n(a)$: fluctuation local, bias global (the A7 split).

Backs the locality verdict and the analytic content of the rate hypothesis on a deterministic-functional-of-i.i.d.-bits toy (ring sites; pair-local tension $T(\{a,b\}) = \text{chord}^2 + \epsilon\,\beta_a\beta_b$; metric (30); positions by global MDS and by a locally-supported quasi-interpolation $\Pi^\dagger$). (1) Tension and metric are exactly edge-local (zero non-incident edges change under a single bit flip). (2) τ_{null} is single-flip-stable ($|\delta\tau_{\text{null}}| \leq 2\epsilon$, $\sim$ machine zero): an ensemble cut, not a per-draw order statistic. (3) The paper's local-sum A_n carries the $O(1/|D_n|)$ McDiarmid bounded difference (empirical slope -0.73, no position-independence assumed). (4) Under one bit flip the $\Pi^\dagger$ fluctuation stays within a bounded neighbourhood (93% within ring-distance 8, zero past 12) while the global MDS embedding spreads it (19%, flat in distance). Confirms fluctuation local / bias global: edge-local metric does not imply local fine-placement, the embedding being globally rigid. Stands in for: the emergence map, by a deterministic functional of i.i.d. bits on ring sites.

Computation 126 · §4.1 · [NUMERICAL]

Forced/free decomposition of the tension functional T.

Backs the structural sharpening after (17), by direct enumeration of $\mathcal{P}(D)$. (1) The admissible magnitudes ((14) zero-locus + (15) monotonicity) form a cone of dimension $2^{|D|} - (|D|+1)$, whose ratio to $2^{|D|}$ rises ($0.50 \to 0.96$ for $|D| = 3 \to 8$): the form is not a finite-parameter family, T_{ex} one point of

it. (2) Several distinct magnitudes (uniform/weighted edge counts, a rank potential) are simultaneously admissible. (3) A monotone reparametrisation preserves every threshold super-level set exactly (downstream reads only the *ordering*), while a different magnitude changes it (the contingent per-configuration content). (4) The forced content (zero-locus; $T(D) \neq 0$, so $\hat{\tau}$ is defined) is invariant across all admissible magnitudes. (5) The per-configuration count is $2^{|D|} = 64$ at $|D| = 6$, of which the ~ 20 observed flavour parameters consume $\sim 31\%$.

Computation 127 · §3.4 · [NUMERICAL]

P2's irreducible primitive is non-balance of v, independent of (14).

Backs the reducibility verdict on P2 (the P2 non-degeneracy item of Ch. 19): almost all of P2 reduces to P1 + (14), the surviving primitive being one non-degeneracy condition, the non-balance of the elementary directions $\left|\sum_{a \in D} v(a)\right| > 0$. (1) A balanced assignment ($\sum v = 0$, three antipodal V_7-pairs) has maximal edge count yet $|\sum v| \approx 10^{-16}$, so (14) secures the count but not the vector bias $T(D)$: non-balance is logically independent of (14), vindicating (16) as a separately-adopted constraint. (2) $G_2 \subset SO(7)$ acts linearly, so $g(\sum v) = \sum g(v(a))$: over 2000 random $SO(7)$ maps the balanced set stays balanced (drift $\sim 10^{-16}$) and norms are preserved, refuting any reduction of non-balance to the Hurwitz/G_2 algebra. (3) Among 2×10^5 random unit assignments none lands within 10^{-3} of the balanced locus (codimension seven at generic balanced points; six on the all-parallel stratum, which at $|D| = 2$ is the whole locus, Computation 152): non-balance is measure-zero but, since PST fixes no measure over the structural datum $v\colon D \to V_7$, not forced here. Computation 152 supplies the measure-free replacement.

Computation 128 · §1.5 · [PROOF]

The minimal KO split $10 = 4 + 6$ is forced given the inputs; the only underived substrate choice is the odd-$|D|$ parity.

Backs the minimal-KO-split verdict (§1.5), which Computation 148 then sharpens from a selection principle to a consequence of the two verified KO-dimensions. With the Lorentzian target $\equiv 2 \pmod 8$ (the Connes value) and the parity-selected substrate KO-6 (Computation 3), the spacetime KO s satisfies $s+6 \equiv 2 \pmod 8$, so $s \equiv 4$: among $s \in \{0, \ldots, 7\}$ the unique solution is $s = 4$ (next $s = 12$), making spacetime KO-4 the arithmetic complement, not an independent input (Computation 4 verifies the spin structure on an already-four-dimensional $M = \mathbb{R} \times S^3$). The minimal total $\equiv 2 \pmod 8$ with $s \geq 1$ is $10 = 4+6$ (next $18 = 12+6$), so the split is forced given the substrate six; minimality is non-vacuous only at the lift $s = 4$ over 12, fixed by the Mosco limit. Varying the internal f flips the split, so the odd-$|D|$ parity that fixes $f = 6$ is the single load-bearing underived choice, not the mod-8 arithmetic.

Computation 129 · §18.1.6 · [SURROGATE]

A spectral gap reaches the relaxation exponent, not the stationary-mean coefficient: the A7 bias residual is immune to mixing rates.

Backs the A7 spectral-gap verdict (§18.1.6). The rate-hypothesis residual binds at the leading coefficient $\eta_{K,R} < d_0^2/4$, not an exponent. (1) A one-parameter family of reversible 2-state chains with fixed spectral gap $g = \frac{1}{2}$ has its stationary mean swept across a 0.8-wide range with the gap identical to machine precision: a relaxation rate does not determine the stationary mean. (2) For $A_n(u)$ with a controlled equidistribution defect ε, the systematic mean deviation is flat in $|D_n|$ (matching the predicted εG) while the statistical fluctuation decays as $|D_n|^{-1/2}$ (std halving per 4× sample). So a covariance-

decay / spectral-gap estimate reaches only the fluctuation exponent; the systematic-mean coefficient is decoupled (the stationary distribution is the kernel eigenvector), and the A7 bias half is foreclosed to mixing / spectral-gap methods. Stands in for: any mixing route to the A7 residual, by two-state chains with a fixed gap; a counterexample, not a model of the cloud.

Computation 130 · §18.1.6 · [SURROGATE]

The A7 bias residual c_u is a cell-volume equidistribution defect; the matched-scaling count does not force it.

Backs the A7 deterministic-frontier reduction (§18.1.6). The quantity

$$c_u \;=\; \frac{|D_n|^{-1}\sum_a |\nabla u(\rho_n(a))|^2}{|K|^{-1}\int_K |\nabla u|^2} \tag{284}$$

is the uniform-weight nodal quadrature defect of the emergence cloud, converging to 1 iff the points cell-volume-equidistribute (not the Delaunay finite-element interpolation error). (1) A deterministic volume-equidistributed cloud gives $c_u \to 1$. (2) A matched-count cloud with a central void in a low-$|\nabla u|^2$ region retains a persistent $|c_u - 1| \approx 0.094$, flat across a $64\times$ count increase and matching the analytic void bias to 10^{-4}: refinement (the A1 count) does not close it. (3) The defect grows monotonically with the void radius. So the A7 bias half is conditional on cell-volume equidistribution of the emergence cloud, the deterministic sibling of the A2 weak-discrepancy (its rate A3 supplied by P3's maximum-entropy sublimation), not on the point count. Stands in for: a non-equidistributed emergence cloud, by a cloud with a void at the same count.

Computation 131 · §18.1.6 · [SURROGATE]

The weak A7 discrepancy is blind to the null-cone vanishing-separation set (codim ≥ 1); only the codim-0 node-density binds.

Backs the A2-core characterisation. c_u is a single-node average, so the Lorentzian vanishing-separation locus (a codim-≥ 1 pair relation) cannot move it. On a 3-cube with $s^2 = -\Delta x_0^2 + \Delta x_1^2 + \Delta x_2^2$: (1) injecting exact null pairs drives the strong Lorentzian minimum separation to 0; (2) the weak defect $|c_u - 1|$ is unchanged by that injection (the paired nodes sit at bulk-typical positions), staying ≈ 0; (3) only a codim-0 central void biases c_u, growing with the void volume. So strong separation fails on the null cone while the weak H^1-tested discrepancy survives, and the A7 residual is purely the codim-0 cell-volume equidistribution of the emergence cloud. Stands in for: the emergence cloud under the signed-tension metric, by 3-cube clouds with injected null pairs.

Computation 132 · §18.1.6 · [SURROGATE]

The tension functional is blind to bulk node density: a void costs nothing in T-terms.

Backs the final foreclosure of the A7 internal-route chain. T reaches physics through three channels, all functions of the distinction data ($v\colon D \to V_7$ and δ), not the placement ρ: the pairwise tensions, the magnitude ordering of $|T(C)|$, and the global direction $\hat{\tau}$. (1) Holding the distinction data fixed, all three channels are identical to machine precision across a uniform and a void placement: T takes no spatial-placement argument. (2) The single-node discrepancy c_u reads ρ, so it sees the void, growing with the void radius. (3) The placement-independent toy T-energy changes by exactly 0 across void radii while c_u rises. So a positive-volume void costs nothing in T-terms, and

the codim-0 cell-volume equidistribution at the core of A7 is forced by no internal lever. Stands in for: the placement ρ, by fixed distinction data placed uniformly and with a void.

Computation 133 · §18.1.6 · [SURROGATE]

A metric-preserving void biases only the flat node-quadrature: a coordinate artefact, not a density axiom.

Backs the coordinate-artefact verdict. The tension data fixes the intrinsic metric only up to diffeomorphism; bulk node density is a distinct diffeomorphism-class invariant it does not pin. Two coordinate charts of the same intrinsic 1D manifold with N geodesically-uniform nodes: a uniform chart, and a void chart $x' = \mathrm{CDF}^{-1}(x)$ with a central density dip and metric $g_2 = n^2$. (1) Every pairwise geodesic distance is identical to machine precision ($\sim 10^{-16}$): the void chart is a genuine diffeomorphic representative (same intrinsic geometry, T-realisable). (2) The coordinate uniform-weight node-quadrature c_u, which the bias half requires to approach 1, is unbiased in the uniform chart but biased by ~ 0.10 in the void chart. (3) The bias grows with the void depth. So the flat c_u is chart-dependent; restoring the void chart's own metric factors makes the defect void-blind (Computation 134), and a void with its co-transforming metric is one gauge orbit carrying no independent density freedom. Stands in for: the coordinate freedom of ρ, by two charts of one one-dimensional manifold.

Computation 134 · §18.1.6 · [SURROGATE]

Computation 133's void bias is a dropped-metric-factor artefact; the genuine A7 residual is the fixed-g rate.

Corrects Computation 133. On its void chart ($x' = \mathrm{CDF}^{-1}(x)$, metric $g_2 = n^2$): (1) the flat node-quadrature $c_u = \mathrm{mean}\,|u'|^2 / \int |u'|^2$ is biased ≈ 0.10; (2) the

metric-covariant c_u (intrinsic gradient $|u'|^2/g_2$, volume $\sqrt{g_2}\,\mathrm{d}x' = n\,\mathrm{d}x'$) equals 1 to machine precision, void-blind, so the flat bias is purely the dropped g^{-1} and $\sqrt{g}$ factors; (3) a genuine fixed-metric density modulation (nodes of density n on $g = 1$, not a diffeomorphism) does bias the covariant c_u. So the metric-preserving void is pure gauge (no bulk-density axiom, no new postulate), and the genuine content is the diffeomorphism-invariant leading-coefficient inequality $\eta_{K,R} < d_0^2/4$ for the emergence-determined ρ_n, the equidistribution condition A3 supplied by P3's maximum-entropy sublimation and corroborated on the i.i.d. surrogate. Stands in for: the same two-chart construction as Computation 133, with the metric factors restored.

Computation 135 · §10.2 · [NUMERICAL]
Born rule via Gleason on $L^2(\mathcal{P}(D), \mu)$.

Verifies Gleason's dimension hypothesis ($\dim H = 2^{|D|} \geq 3$) and demonstrates the Born form $P(C) = |\psi(C)|^2\mu(C)$ on the substrate Hilbert space for $|D| = 2, 3, 4$: total probability one, frame-function additivity in an arbitrary orthonormal basis, and non-contextuality of shared projectors. Confirms the $|\psi|^2$ rule is the Gleason frame function of the pure state, not an added postulate.

Computation 136 · Ch. 13 · [PROOF]
Gauge anomaly cancellation and the $N_c = 3$ condition.

Evaluates all six triangle-anomaly coefficients for one Standard-Model generation as exact functions of the colour rank N_c. All vanish at $N_c = 3$; the $[\mathrm{U}(1)_Y]^3$ coefficient is $-2N_c+6$ and the $[\mathrm{SU}(2)]^2\,\mathrm{U}(1)_Y$ coefficient is $(N_c-3)/6$, so anomaly freedom (given the SM hypercharges) forces $N_c = 3$, the dimension of the $M_3(\mathbb{C})$ colour factor. A consistency check, not a derivation of the hypercharges.

Computation 137 · §16.5 · [PROOF]

Equation of state $w = -1$ from steady-state non-dilution.

Two exact routes to $w = -1$: (A) the FRW continuity equation $\dot{\rho}+3H(\rho+p) = 0$ with the non-diluting substrate density $\dot{\rho} = 0$ forces $\rho + p = 0$; (B) a metric-proportional stress tensor $T_{\mu\nu} = \rho_\Lambda g_{\mu\nu}$ gives $T^\mu{}_\nu = \rho_\Lambda \delta^\mu{}_\nu$, matching a perfect fluid at $p = -\rho$. The load-bearing input is non-dilution; $w = -1$ follows from it.

Computation 138 · §14.3 · [EMPIRICAL]

CKM consistency with the S_3-leading structural-scope theorem.

Checks that the measured CKM matrix (PDG) is consistent with identity plus a small $T(C)$-contingent perturbation, as the S_3-leading theorem requires: $\|V-I\|_2 \approx 0.25$ is Cabibbo-sized ($O(\lambda)$, not $O(1)$), every off-diagonal is $\leq |V_{us}|$, and the magnitudes are unitary to sub-percent. A falsifiable consistency bound (the values themselves are out of scope).

Computation 139 · Ch. 15 · [NUMERICAL]

The d^{-6} Casimir exponent derived from the projection kernel.

Derives the d^{-6} pressure-correction exponent directly from the mode form factor rather than by dimensional analysis: for any smooth even kernel $g(kd_0)$ ($g(0) = 1$, $g'(0) = 0$ by parity), the leading correction is $\delta P = -C\, a_2 d_0^2\, \Gamma(6)\zeta(6)/(2d)^6 \sim d_0^2 d^{-6}$. Numerically confirmed for four smooth even kernels (fitted exponent -6.00 to $< 1\%$; analytic coefficient matched to $< 0.2\%$); the exponent is kernel-*independent* (a non-smooth cusp kernel gives -5 instead), while the coefficient a_2 is kernel-*specific* (varies $4\times$), so the d^{-6} signature is generic to any smooth even single-scale cutoff and only $\xi = 90/\pi^2$

is PST-specific.

Computation 140 · §4.1 · [PROOF]

Robustness of the emergence chain to the form of $T(C)$.

Formalises and verifies that the downstream structural derivations read $T(C)$ only through the descending order of $|T(C)|$ (the instantiation sequence) and the global bias $\hat{\tau} = T(D)/|T(D)|$. For any strictly increasing reparametrisation ϕ ($\phi(0) = 0$) of the magnitudes with $\hat{\tau}$ fixed, at matched threshold $\tau' = \phi(\tau)$ the instantiation sequence, the per-configuration phase $\text{sign}(|T(C)| - \tau)$, and $\hat{\tau}$ are all invariant (300 random trials); a non-order-preserving reassignment changes the sequence. So the residual freedom in $T(C)$ beyond the ordinal content and $\hat{\tau}$ is inert. Does not derive a canonical $T(C)$ (open, Computation 126); it fixes the exact robustness criterion.

Computation 141 · Ch. 3 · [NUMERICAL]

The directional marginal over v is forced by G_2; only independence is adopted.

Sharpens Computation 127, which closed with "PST fixes no measure over v" and concluded that non-balance is adopted rather than derived. A measure over v factors into a directional marginal on S^6 and a radial one, and neither is what remains adopted. Builds G_2 itself as the derivation algebra of the octonion cross product (dim = 14; Computation 127 used random SO(7) maps as a proxy), and confirms with genuine automorphisms that the balanced locus is G_2-invariant, so the Hurwitz/G_2 *exclusion* route is refuted. But the orbit through a unit vector is 6-dimensional: G_2 is transitive on S^6, which therefore carries a *unique* G_2-invariant probability measure, so the directional marginal is derived in the same way the Bernoulli measure is. The radial marginal is then shown irrelevant: with directions drawn from that measure

independently across properties, no balanced assignment occurs in 200 000 trials at $n = 2, 3, 8$ for unit, uniform, exponential, heavy-tailed or six-decade radial laws. The one exception is degenerate and self-closing: if $v(a) = 0$ is permitted with probability p then $P(\text{balanced}) = p^n$ exactly (measured against prediction to three figures), which vanishes as $|D|$ grows for every $p < 1$. What is load-bearing is independence: a law coordinating the values across properties makes balance certain, an independent one makes it null. So what is still adopted is not a measure but independence, the structural claim that nothing at substrate level coordinates the values, which is exactly the ungroundedness argument. It supplies the genericity argument's missing premise rather than replacing it.

Computation 142 · §10.5 · [NUMERICAL]

Tensor factorisation over a bipartition, and no-signalling.

Backs the composite-system section. PST's state space is $L^2(\mathcal{P}(D), \mu)$ with μ the Bernoulli product measure of §2.5; because μ is a product over the index set D, it factors over any bipartition $D = D_A \sqcup D_B$, and the Hilbert space factors with it. *Findings*: (1) $\beta(C) = (C \cap D_A, C \cap D_B)$ is a bijection $\mathcal{P}(D) \to \mathcal{P}(D_A) \times \mathcal{P}(D_B)$, checked exhaustively at $|D| = 12$ over four splits, 4096 configurations each. (2) The measure factors to 0.000, with marginals taken from the measure itself rather than assumed; a *non-product control* (weight $\propto 1 + |C|$, permutation-invariant but not a product) fails the same test at 9.9% of its largest weight, so the tensor structure rests on the Bernoulli measure specifically and not on permutation invariance alone. (3) The induced map preserves inner products (10^{-17}) and elementary tensors reproduce the Kronecker product (7×10^{-18}). (4) Over 200 random correlated states and random B settings, Tr_B of the ensemble-averaged post-measurement state is invariant to 6×10^{-16}: *no-signalling*. (5) The same invariance holds on the singlet, whose CHSH value is confirmed at $2.828427 = 2\sqrt{2}$; correlations beating the classical bound still signal nothing. Conclusion: composite systems,

the reduced state and no-signalling are all consequences of the canonical measure, not additional assumptions.

Computation 143 · §10.3.2.5 · [NUMERICAL]
Born compatibility of the threshold-crossing rule.

Asks whether the crossing rule of (98) selects outcomes with the Born frequencies that §10.2 derives independently. For a fixed state the weights w_i are fixed, so the randomness is in the configuration; the object is $\mathbb{P}_C(w_i > \tau_i(C))$. Result: the rule reproduces Born *exactly*, $\mathbb{P}(\tau < w_i) = w_i$, precisely when the tension variable is uniform in its quantile; verified on three states to sampling precision and shown to fail on the same states for a Beta$(2, 2)$ control. Also maps the uniqueness/Born trade-off along the coupling axis: raising $\kappa\lambda^2$ drives the probability of two simultaneous crossings monotonically to zero (the hypothesis clause (U) rests on), while pushing further orders the outcomes by margin and makes the survivor deterministic. A control placing the largest weight on the largest $|o_i|$ rules out a labelling artefact.

Computation 144 · §19.9 · [NUMERICAL]
Observable trace of a slit-plane crossing: the duality fork.

Tests whether the intermediate-coupling regime predicts anything standard quantum mechanics does not, in a way that does not depend on the details of the threshold law. If a crossing at the slits physically selects a path, only the fraction $1 - p$ of runs without one shows fringes, giving $V^2 + D^2 = 1 - [1 - (1 - p)^2]\gamma^2$, a sub-duality deficit maximal as $\gamma \to 1$, i.e. at *weak* coupling; verified against the numerical maximum to 10^{-6}, monotone in p and zero only at $p = 0$. A duality test good to ϵ excludes $p > 1 - \sqrt{1 - \epsilon}$. With p near 0.5 from Computation 143 the predicted deficit is gross, so the selecting reading is untenable and crossings must be redundant with the inherited

decoherence.

Computation 145 · §10.3.2.5 · [NUMERICAL]

Is the Bernoulli tension uniform in its quantile?

Supplies the condition Computation 143 identifies, by exact enumeration of all $2^{|D|}$ configurations for the scalar exemplar T_{ex} of (17). Raw, the tension is far from uniform (KS 0.46–0.51 for every distinction graph tried, at $|D| = 14$). But downstream reads only the *ordering* of magnitudes and Computation 140 verifies monotone-reparametrisation invariance, so the probability-integral transform is available and makes the variable uniform by construction. The residual obstruction is ties, which no order-preserving map can separate and which floor the achievable KS at half the largest tie mass. Every measured value sits at its own floor. For generic Erdős–Rényi structures the floor is 0.034 at $|D| = 14$ and decays as $|D|^{-1.51}$ ($R^2 = 0.998$), reaching 8.5×10^{-4} at $|D| = 256$; for regular graphs it is large and uniformity is unreachable, but those are the graphs (17) already excludes. *Superseded*: Computation 180 shows both halves of that last clause wrong (§10.3.2.5), and Computation 187 corrects the exponent.

Computation 146 · §10.3 · [NUMERICAL]

What the substrate layer contributes that decoherence cannot.

Fixes the division of labour between the three closures. Evolves a system–apparatus–environment state unitarily and confirms both halves of the standard result: the off-diagonal elements of the reduced state decay to zero, so interference is destroyed and the pointer basis is selected; and the global state stays pure ($\operatorname{Tr}\rho^2 = 1$ to machine precision at every environment dimension), so nothing has become definite. The decohered reduced state is then compared entry by entry with a proper mixture built by hand from definite outcomes:

identical, and indistinguishable under 4000 random Hermitian observables to 0. Global purity separates the two cases perfectly (1 versus 0.58), so the difference is real but not locally accessible. Hence Gleason supplies the probabilities, einselection the basis, and the threshold crossing the definiteness; and Computation 144's null visibility result is what the formalism requires rather than evidence of an idle layer.

Computation 147 · §16.10 · [NUMERICAL]
The reservoir question under the equation of state.

Tests whether ongoing sublimation can exhaust the configuration space. The appeal to an infinite $\mathcal{S}$ does not settle it, since μ is a probability measure with total mass 1 however large D is. Two readings are separated: a *depleting* one, in which a crossed configuration leaves the ensemble, and a *modal* one, in which it is re-labelled. Depletion removes configurations in tension order and biases the survivors away from equidistribution, measured as KS drift in the tension quantile over all 2^{16} configurations; the drift is monotone but slow, first exceeding three times the atom floor of Computation 145 only near 80% sublimated, so this is a weak argument and is reported as one. The decisive exclusion is independent and carries no tolerance: §16.5 already requires j_τ to be a pre-spatial saturation density rather than a per-proper-volume creation rate, since the latter gives $w = 0$, and a draining stock is precisely such a flux.

Computation 148 · §1.7 · [PROOF]
Are the three extremal selection principles independent?

No: they reduce to one. Enumerates the Hurwitz classification ($\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}$; dimensions $1, 2, 4, 8$) and shows that *maximal directional structure* and *maximal-tensor envelope* are a single instruction evaluated on two constructions: argmax over vector parts (dimensions $0, 1, 3, 7$) returns $\mathrm{Im}(\mathbb{O}) = V_7$, argmax over ten-

sor products of distinct members returns $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ at dimension 64, and one function implements both. $\mathbb{R}$ being the tensor unit is why the envelope is a triple product, a consequence rather than a choice. Separately, *minimal split* is shown not to be a selection: with KO-6 verified (Computation 3) and KO-4 verified (Computation 4), the admissible (n, k) with $n + k \equiv 2 \pmod 8$ collapses to the single pair $(4, 6)$, so nothing is minimised; had the summands been free, minimality would have returned $(0, 2)$, which the theory does not use. The reduction does not remove the odd-$|D|$ parity restriction Computation 3 carries.

Computation 149 · §17.6 · [PROOF]

Why one instantiation rather than many.

Tests whether P3's necessity of crossing commits PST to a multiverse. It does not: each route to a second instantiation is a motion along a symmetry the theory already carries. A different direction is a G_2 rotation, and transitivity on S^6 is confirmed by the stabiliser dimension count $14 - 8 = 6 = \dim S^6$, so S^6/G_2 is a point; 2000 sampled directions are transported onto a common reference to 10^{-10}, confirming the orbit is exhaustive. A different realisation map is diffeomorphism invariance: over 500 relabellings the metric's value multiset under (30) is permuted and not altered, to 10^{-12}. A disjoint region factors rather than detaches, the measure factorisation holding to machine precision. The quotient by the three has exactly one element. Stated as a structural result only: it does not bear on the eternal-inflation multiverse.

Computation 150 · §18.3 · [NUMERICAL]

Scope of the L_comm uniform bound, and the $m(k)/N$ mechanism.

Re-measures $G(k, N) := |\mathrm{gap}(k, N)| \cdot N$ over $k \leq 40$, $N \in [12, 40]$ (513 pairs), against Computation 65's window $k \in [1, 16]$. Three findings. (i) The bound is

not uniform in (k, N): $G = 22.0$ at $(k, N) = (39, 22)$, so the previously printed "≤ 5.3 uniformly in (k, N)" is refuted. (Computation 153 shows 22.0 was a property of this sweep's window rather than a maximum: G is unbounded off the operative regime, reaching 38.1 simply by looking further.) (ii) Every violation lies at $k > k_D = N/2$, which the Walsh cutoff excludes; on the operative regime $k \leq k_D$ the sharp constant is 3.674 over 195 points (median 0.95), attained at the *smallest* substrate size $D = 36$. Computation 65's reported maximum 5.249 at $(15, 20)$ has $k_D = 10$ and so lies outside the regime the theory uses. (iii) The controlling parameter is the bridge degree over the truncation, $m(k)/N$: bucketed, $\max G$ is 2.6 for $m(k)/N \leq 0.25$ and 22.0 at $m(k)/N = 0.5$. Since $m(k) \leq 2\sqrt{k}$ (verified for every $k \leq 2000$), the matched scaling gives $m(k_D)/N(D) \leq D^{-1/4}$, so the deficit *vanishes* rather than staying $O(1)$, correcting an earlier step that bounded $m(k) \leq \sqrt{D}$, which is the bound on k rather than on m. An explicit analytical constant remains open; this supplies the variable, the mechanism and the rate, not the value. *Settled by Computations 213 and 214*: the limit is in closed form and the sharp value is derived at $(6, 12)$.

Computation 151 · §19.9.2 · [PROOF]

No-signalling is bipartition-universal, hence ρ-independent.

Closes the first of the two items left open in §19.9.2. The emergence map ρ is not unique, so $D_A = \rho^{-1}(U_A)$ is fixed only relative to a choice, but the theorem it feeds is universally quantified over bipartitions, and no step in it (factorisation, Born rule, $\sum_j Q_j = \mathbf{1}_B$, cyclicity on B) mentions ρ, a metric or a region. Every ρ induces some bipartition and none is exempt, so the conclusion is ρ-independent even though the naming is not. Checked *exhaustively*: the bijection β and the factorisation of μ hold on all $2^{10} - 2 = 1022$ non-trivial bipartitions at $|D| = 10$, to exact arithmetic, across every split size 1 to 9; the A marginal is invariant under every B-side basis to 4×10^{-16}. A permutation-invariant but non-product measure fails factorisation

by fourteen orders of magnitude, so the result is carried by P1's Bernoulli measure. Once factorisation holds the no-signalling step is an identity; the content is that factorisation holds for *every* bipartition. *The residue moved on*: Computation 156 shows the *compactness* half is a corollary of (30) (the distance is $|T-\tau_{\text{null}}|^{1/4}$, strictly monotone, so compact in M and tight in T are one condition), and that being a device is an open condition with sharp margin $m/2$; what is genuinely left is why T clusters, which is $T(C)$-contingent.

Computation 152 · §3.4 · [PROOF]

Non-balance is an open condition, and symmetry does not force it.

Replaces the circular measure-zero defence, which invoked a measure in the same breath as conceding PST fixes none over v. (i) The balanced locus B is closed of codimension seven, so its complement is open and dense and the distance to balance is exactly $\left|\sum_a v(a)\right|/\sqrt{|D|}$: stability, not probability, and no measure used. Read at *generic* balanced points: codimension seven for $|D| \geq 3$, six on the all-parallel stratum, which at $|D| = 2$ is all of B. (ii) The assignments fixed by the relabelling group are the constant ones, and those are maximally *un*balanced, $\left|\sum_a v(a)\right| = |D|\,|v|$, which answers the objection that an unstructured substrate should cancel. (iii) **But symmetry does not force non-balance**: $v(a) = \epsilon_a u$ with shared u and i.i.d. signs is invariant under diagonal G_2, all relabellings and per-property complementation, has unit directions, and still balances with probability $\binom{|D|}{|D|/2}2^{-|D|} \approx 0.31$ at $|D| = 6$ (measured 0.3124). The residue is a coordination premise, weaker than independence: one property anywhere in D whose conditional has no atom at the cancelling value suffices. (iv) *Scope, stated not hidden*: if directions are quantised to $\pm e_1 \ldots \pm e_7$ balance is not null, exactly $1/14$ at $|D| = 2$, so the argument needs P1's continuum of directions.

Computation 153 · §18.3 · [NUMERICAL]

The L_comm rate is exactly $\Theta(1/\sqrt{D})$, and why it looked better.

Closes a question Computation 150 deliberately left open. Its seventeen matched-scaling points looked like they trended down, which would have made $\gamma_D = o(1/\sqrt{D})$, better than the book claims; Computation 150 refused to claim it. The refusal was right. (i) At *fixed* k, G converges to a nonzero constant as N grows even while $m(k)/N \to 0$: at $k = 1$ it runs $1.0161, 1.0128, \ldots, 1.0057$ over $N = 16$ to 44, monotone towards about 1.00, with $k = 2$ and $k = 4$ settling near 0.87 and 0.95. So gap $\sim C/N$ with $C = O(1)$ and $\gamma_D = \Theta(1/\sqrt{D})$; the rate is tight, and the apparent improvement was never there. No conflict with Computation 150: a relative deficit $O(m/N)$ on the norms and an $O(1/N)$ gap in their ratio are different statements. (ii) The apparent trend is sawtooth phase in the $m(k)$ stratum, and Computation 150's window spans $m = 4$ to 8. (iii) The controlling variable sharpens from $m(k)/N$ to the truncation headroom $s = N - 2(m + 1)$ (equation (281)); $m(k)/N$ is a proxy because $m(k)/N \to 1/2$ is the same event as $s \to 0$. Off the operative regime G is *unbounded* as s closes (at $N = 40$, $s = -2$ it runs 3.9 to 38.1 over ten consecutive k) so Computation 150's 22.0 was a property of that sweep's window, not a maximum, and no maximum should be quoted. The analytical constant remains open.

Computation 154 · §8.3.3 · [PROOF]

The derivative expansion does not fail at an astrophysical horizon.

Evaluates the book's own validity condition, the higher-curvature terms of equation (49) are suppressed by $(k\, d_0)^2$ and negligible at scales $\gg d_0$, on the Schwarzschild geometry, rather than assuming the strong-field case lies outside it. (i) Symbolically, from the metric: $R_{\mu\nu} = 0$, $R = 0$, and $K = R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} = 12r_s^2/r^6$ is the only non-vanishing invariant. (ii) With $k = K^{1/4}$,

the suppression factor at the horizon is $\sqrt{12}\,(d_0/r_s)^2 = 2.1 \times 10^{-24}$ for one solar mass, 4.8×10^{-27} for Cygnus X-1, 1.1×10^{-37} for Sagittarius A* and 5.0×10^{-44} for M87*: a horizon is strong *potential*, not strong curvature, and the covariant expansion sees only the latter. (iii) The expansion first fails at $r_{\mathrm{break}} = (12^{1/4} d_0)^{2/3} r_s^{1/3}$, which is 38 μm for a solar mass (1.3×10^{-8} of r_s, enclosing 2.1×10^{-24} of the interior) so the uncovered region is the singularity's neighbourhood, not the horizon. (iv) Mass floor 2.9×10^{18} kg, where $r_s = 12^{1/4} d_0 \approx 4.3$ nm; with the grain ℓ_* as cutoff instead of d_0 it is 1.9×10^8 kg, so the conclusion is insensitive to which PST length controls the expansion. (v) Every term of the field equations and of the Noether-charge integrands of the named corrections $\mathcal{O}(R_{\mu\nu}R^{\mu\nu}, R^2)$ carries a factor R or $R_{\mu\nu}$, so the black hole solves the *truncated* equations exactly and its Wald entropy is $A/4G$ with no $(d_0/r_s)^2$ correction; the derived Λ costs a relative 10^{-119}. (vi) The unnamed $R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$ is Gauss–Bonnet up to those two, and its Noether charge is $64\pi^2\lambda$, mass-independent: an additive $O(1)$ against $S/k_B = 1.05 \times 10^{77}$, not an area term. The formula is calibrated in the same script by checking that it returns $A/4G$ from equation (50). Does *not* establish the global causal structure (event versus apparent horizon turns on the censored interior, exactly as in general relativity), does not make the substrate-side horizon identification a theorem, and says nothing about the statistical entropy. *Scope, as first run*: on the 2.3 nm Casimir bound of its day; Computation 216 re-evaluates on $d_0 = 2.2$ am (37 pm, 1.9×10^{-42} of the volume, a floor of 2.8×10^9 kg), every conclusion surviving by wider margins.

Computation 155 · §7.4 · [PROOF]

The tensor projection $\Pi_{\mu\nu}$: structure, mechanism, and what is left.

Three results on the open tensor generalisation of the stress-energy projection. (i) *Scope of the scalar form, symbolically on Friedmann–Robertson–Walker.* The Bianchi identity applied to $T_{\mu\nu} = \Pi g_{\mu\nu}$ gives $\nabla^\mu T_{\mu\nu} = \partial_\nu \Pi$, so conservation forces Π constant; and a perfect fluid is proportional to $g_{\mu\nu}$ iff $\rho + p = 0$.

The scalar form is therefore exactly and only the cosmological constant, not "the leading-order approximation valid for isotropic configurations": isotropic dust and radiation are already outside it. (ii) *The obvious completion is singular on the light cone.* The unit dyad $\hat{e}_\mu \hat{e}_\nu$, $\hat{e} = \Delta/\sqrt{|\Delta \cdot \Delta|}$, diverges as $\Delta \cdot \Delta \to 0$, which is where τ_{null} sits by (30); measured, the summed per-pair value swings 56% over clouds of 200 to 800 points and does not settle. Any completion normalising by the interval fails the same way. (iii) *The separation dyad works, and its trace is the tension.* With $\Pi_{\mu\nu}$ as in (40), $g^{\mu\nu}\Delta_\mu\Delta_\nu = s_{ab}$ exactly (checked pair by pair, error 0), so the scalar Π returns as the *trace* rather than as a factor on the metric. For any invariant-only weight on a rotationally isotropic ensemble the result is forced into perfect-fluid form, residual falling $1.5 \times 10^{-1} \to 4.8 \times 10^{-2}$ as the cloud grows $200 \to 1600$. ρ is the timelike-pair moment and p the spacelike-pair moment, split by $T(\{a, b\})$ against τ_{null}, the same cut that fixes the signature, so $w = p/\rho$ is the balance of the two tension budgets. Solving $\rho + p = 0$ in the tested family gives $\lambda^* = 0.8699$, where $\|\Pi - \lambda g\|/\|\Pi\| = 3.7 \times 10^{-2}$ equals the decomposition residual: at that point the tensor *is* the scalar form, so the cosmological constant is one point of a one-parameter family. *Superseded in part by Computation 158*, which supplies what this computation left open and corrects one of its conclusions. The weight W is *not* adopted: trace consistency between (31) and (40), with (30) giving $T - \tau_{\text{null}} = -s|s|$, forces $W(s) = -|s|$. And the reachable range reported here, $w \in (-1, +0.200]$ on the physical branch, from which this computation concluded that radiation at $+1/3$ was out of reach, was a property of its two-class exponential family rather than of the construction: the derived weight gives $w = 1/3$ exactly. The λ^* result stands as stated, for the family it was measured on. Still open from this computation's list: conservation is checked flat rather than on a curved background, and the tensor is fixed only up to a divergence-free term.

Computation 156 · §19.9.2 · [PROOF]

What would make a property subset a device, and what would not.

Takes apart the one item §19.9.2 still carried after Computation 151: that a laboratory device occupies a compact region "is inherited rather than derived". Three pieces, one of them genuinely open. (i) *The compactness half is a corollary.* (30) gives $|s^2| = |T - \tau_{\text{null}}|^{1/2}$ for a spacelike pair, so $d(\rho(a), \rho(b)) = |T(\{a, b\}) - \tau_{\text{null}}|^{1/4}$, strictly monotone: "compact in M" and "tight in T" are one condition written twice, and a planted subset is recovered from the tension matrix alone (round trip below 10^{-9}). (ii) *Stability, sharp and measure-free.* With $m = \text{sep} - \text{diam}$, every perturbation of the induced distances below $m/2$ preserves the property and $m/2$ is sharp: adversarially, $0.99\, m/2$ intact, $1.01\, m/2$ broken (measured diam $= 0.0734$, sep $= 0.2152$, $m = 0.1417$). A completeness lemma makes the search exhaustive: diam $<$ sep forces D_A to be exactly the $|D_A|$ nearest neighbours of each of its members, confirmed against brute force over all subsets at $|D| = 14$ (both find five). (iii) *Cardinality is what makes the criterion bite.* Bounded diameter with separation is on its own far too weak; every nearest-neighbour pair satisfies it, rate 1.000. Over 200 replicas of 60 properties the accident rate runs 1.000, 0.965, 0.300, 0.040 at sizes two to five and 0.000 from six up, so a twelve-property device is not an accident of an unclustered configuration. *What is left*: why T clusters at all, which is $T(C)$-contingent rather than an open derivation, plus the bound-state physics the projected field theory already inherits. *Caveat, the same one Computation 152 raised*: the accident rates use a uniform null and PST fixes no measure over configurations, so they are indicative; the margin result is the measure-free part.

Computation 157 · §19.2.1 · [PROOF]

The L_comm bound reduced to ladder commutators, and solved at $k = 1$.

Attacks the structure of the L_comm constant rather than sweeping for it. (i) *The Dirac operator drops out.* T_{w^m} preserves $a - b$, so $[J_z, T] = 0$ exactly for every bridge, and the $\sigma_\pm$ algebra block-diagonalises the rest: $\|[D_\alpha, T \otimes I_2]\| = \max(\|[J_+, T]\|, \|[J_-, T]\|)$ (equation (273)), the two being equal by the $z_0 \leftrightarrow z_1$ symmetry. Checked to 10^{-14}. (ii) *The ladder commutator is a shift*: $[J_+, T_{w^m}] = m\, S_{m+1,m-1}$ exactly, to 10^{-16}. Bridge and commutator are then each a sum of two shifts differing by $(1, 1)$, hence bidiagonal along the chains $(a, b) \mapsto (a + 1, b + 1)$; the resulting $O(N^2)$ evaluator reproduces the brute-force singular value to 10^{-11}. (iii) *Closed form at $k = 1$*, where $m = 1$ and $\alpha = 0$ so both operators are single shifts: $\|T\|^2 = (N^2/4)/((N + 1)(N + 2))$ at $a = b$ and $\|[J_+, T]\|^2 = (N - 1)N/((N + 1)(N + 2))$ at $b = 0$, giving $\text{ratio}(1, N)^2 = 4(N - 1)/N$ (equation (274)) and $C(1) = 1$ exactly, derived, not fitted, and matching to 10^{-12} up to $N = 50$. (iv) *The sharp constant re-checked four times further out*: Computation 150's sweep tested the operative regime only to $k = 20$; at $k \leq 80$ the maximum of G is still 3.6736 at $(k, N) = (6, 12)$, the smallest admissible N for $k = 6$, with per-k maxima falling to 1.5–2.2 by $k = 50$–80. (v) *The limit constants are a different number*: near $N = 250$, $C(k)$ is 1.0011, 0.83, 0.10, 0.78, 0.79 at $k = 1, 2, 3, 4, 6$, the first matching $1 + 1/(4N)$ from the closed form. *Still open*: the value. A closed form for general m needs the discrete maximum over a, which is what generates the sawtooth.

Computation 158 · §7.4.4 · [PROOF]

The weight profile W is not free, and what it forces.

Closes the weight-profile question §7.4 handed on. *The derivation*: (31) projects the tension itself, the trace of (40) is $g^{\mu\nu}\Delta_\mu\Delta_\nu = s$, and (30) gives $T(\{a,b\}) - \tau_{\text{null}} = -s|s|$ identically ((42); checked for timelike, null and spacelike pairs). Requiring the two projections to agree on the trace, pair by pair, is one equation in one unknown and gives $W(s) = -|s|$ ((43)) with no free function and no free scalar beyond the $8\pi G$ unit conversion. *One addition, from which nothing is read*: without a range restriction the long pairs dominate and w drifts with the sampling window (spread > 0.05 over windows 0.5 to 4), which a density must not do; with the coherence scale as cutoff the answer converges, identically for four cutoff shapes (exponential, hard, Gaussian, algebraic) at two scales differing fourfold. *What it forces*: the trace of $\Pi_{\mu\nu}$ *is* $\sum(T - \tau_{\text{null}})$, and it vanishes, the null threshold being the balance point of the pair-tension distribution; $(\rho - 3p)/\rho$ falls $-1.9 \times 10^{-1} \to -5 \times 10^{-6}$ as the window grows, so $w = 1/3$ exactly ((44)). Independent of the configuration: uniform, radially graded, Gaussian in space, Gaussian in space and time, and ten-clump clustered all give 0.3333 to 0.3338. The pair-tension projection is conformal matter. *This also overturns Computation 155's ceiling*: that computation's two-class family reached only $w \leq +0.200$ and flagged radiation as possibly unreachable with the caveat that the ceiling might be an artefact of the exponential profile. It was. *Open*: dust, which breaks the forced tracelessness and needs a mass scale the pair-separation structure does not carry; Λ is not on that list, being the separate vacuum term with $w = -1$ (Computation 155, §16.5). *Scope, from Computation 199*: the trace vanishes in the large-window balanced limit; out of balance it is the balance scalar, and the identification concerns the conformal part of the sum.

Computation 159 · §7.4.6 · [PROOF]

Dust is not in this sector, and a photon is why.

Settles the dust question Computation 158 left open, in the negative, which locates the rest mass rather than leaving it vague. (i) *The worldline sector is exactly kinetic.* Real matter is worldlines, not scattered points; restricted to intra-worldline pairs the projection gives $w = \beta^2/3$ ((47)) exactly, to 5×10^{-15} over nine speeds, with $p = 0$ at $\beta = 0$ and $1/3$ as $\beta \to 1$. Nothing is put in: it follows from the forced weight (43). (ii) *A photon contributes nothing from its own worldline.* A null worldline has $s = 0$ along it, so $W(s) = -|s|$ vanishes identically; measured, its summed weight is fourteen orders below the $\beta = 0$ case and consistent with round-off. Radiation's $1/3$ is therefore forced by the weight vanishing on null separations, the separation τ_{null} defines, so the photon enters the stress tensor through the same object that fixes the signature. (iii) *Mixture law* $w = [1-f(1-\beta^2)]/3$ with f the worldline share, to 5×10^{-5}: dust needs $f \to 1$ and $\beta \to 0$. (iv) *f does not reach one, structurally.* $|s| \lesssim d_0^2$ is a neighbourhood of the light cone, not a compact region: in a box of half-width 40, 98% of surviving pairs are separated spatially by more than seven correlation lengths, median 31, so the near-null background never dilutes and f stays at 0.01 to 0.05 over fifty times the correlation length. (v) *The bond kernel cannot rescue it*: B is normalised, $\int B \, \mathrm{d}^4x = 1$, so the volume-averaged density is identically the integrated sum for any B (measured 0.33329 against 0.33337); it decides where the stress sits, not the equation of state. *Conclusion*: the projection is strictly conformal for homogeneous matter, dust is excluded rather than unfound, and the rest mass enters through the matter action rather than the pair geometry.

Computation 160 · §8.3.4 · [PROOF]

The photon: what τ_{null} fixes, and what is still missing.

Collects four facts about light that are usually presented as unrelated and shows one object behind them. (i) *Two polarisations*: the physical states of a vector excitation are $\{\epsilon : k \cdot \epsilon = 0\}$ modulo $\epsilon \sim \epsilon + \lambda k$; the transverse subspace is three-dimensional for every k, and the gauge direction is removable only when k lies in it, which is $k \cdot k = 0$. Null gives two states, timelike gives three, over 200 random momenta each. The counting is ordinary gauge theory; the input, that the photon's momentum is null, is PST's, through the unbroken U(1) and (30). (ii) *Equation of state*: a null worldline has $s = 0$, so the derived weight vanishes identically on it and $w = 1/3$ is forced (Computation 159). (iii) *Stefan–Boltzmann*: $p = \rho/3$ makes the first law integrable only for $T \propto \rho^{1/4}$, so $\rho \propto T^4$, solved symbolically; no quantum input is used, and PST reaches it because it derives the $w = 1/3$ the argument needs. (iv) *The coefficient too*: $a = g(\pi^2/30)k_B^4/(\hbar^3 c^3)$ with $g = 2$ gives $7.5657 \times 10^{-16}\ \mathrm{J\,m^{-3}\,K^{-4}}$, the measured value to eight figures, where $g = 3$ would miss by half. *Missing, and now supplied*: this computation said the Planck spectral distribution was unassembled and that the quantisation into countable quanta was absent outright. Computation 161 assembles the spectrum, and the second half of that claim was wrong: §12.1 has the Fock space by canonical quantisation, inherited rather than derived. The photoelectric effect is genuinely absent and stays so.

Computation 161 · §8.3.5 · [PROOF]

The Planck spectral distribution, assembled.

Writes down the derivation Computation 160 said was missing, and separates what PST forces from what it inherits. *Four inputs are forced*: the linear dispersion, from masslessness at τ_{null}; $g = 2$, from the null condition; *bosonic*

occupancy, because the photon is an inner fluctuation of even Clifford grade (Computation 108) and even-order composites are bosonic (§11.3); and $\mu = 0$, because the conserved U(1) charge is the winding number of $\pi_1(\mathcal{V})$ and *not* the occupation number, so nothing conserves photon number. *One is inherited*: the ladder $E_n = n\hbar\omega$ comes from canonical quantisation of (123) in the ordinary way (§12.1); PST supplies the arena, not the ladder. *Checked*: the bosonic, fermionic and classical branches of the mean occupation, symbolically, the last being the ultraviolet catastrophe; the density of states $g\omega^2/2\pi^2c^3$ against a brute-force count of box modes, the residual shrinking as a surface term; (69) against the textbook angular-frequency, frequency and wavelength forms; the integral $\int_0^\infty u\,\mathrm{d}\omega$ by term-by-term expansion and resummation to $6\zeta(4)$, cross-checked by quadrature to 2×10^{-13}, returning *exactly* the radiation constant of Computation 160 ($g = 3$ would be half as large again); Wien's displacement law in both conventions, $b = 2.8977720 \times 10^{-3}$ m K against CODATA; and the Rayleigh–Jeans and Wien limits. *The falsifiable part*: a conserved photon number would give $\mu \neq 0$, and the measured microwave-background spectrum bounds the distortion at $|\mu| < 9 \times 10^{-5}$ [209], so the structural $\mu = 0$ is confirmed to a part in 10^4. *Still open*: the photoelectric effect, which needs the interaction vertex rather than the free spectrum.

Computation 162 · §7.4 · [PROOF]

Conservation of $\Pi_{\mu\nu}$ on a curved background: bounded, not closed.

Backlog M3, and the first of the two items §7.4 leaves open. *The covariant reading*: a pair has no separation vector on a curved manifold, so Δ_μ becomes Synge's world function derivative σ_μ with $s = \sigma^\mu\sigma_\mu = 2\sigma$. *Two exact identities*, that one and its gradient $\sigma^\mu\sigma_{\mu\nu} = \sigma_\nu$, collapse the per-pair divergence to (45), so every curvature effect enters through the single scalar $\Box\sigma$. *In closed form on constant curvature*, by the Laplace–Beltrami operator in geodesic polar coordinates, $\Box\sigma = 1 + (n-1)(d/L)\cot(d/L)$ on the sphere and coth on the hyperboloid, both tending to n as $L \to \infty$ and both reproducing the

coincidence expansion $n - \frac{1}{3}R_{\alpha\beta}\sigma^\alpha\sigma^\beta$. *The shift*: with the derived weight the relative correction is $-\frac{1}{7}(d/L)^2$, second order as it must be. *The bound*: $|s| \lesssim d_0^2$ gives $d \lesssim d_0$, so the violation is 6×10^{-25} at a solar-mass horizon, 6×10^{-27} at a neutron star's surface, 9×10^{-41} at the Earth's and 6×10^{-70} cosmologically, the same $(kd_0)^2$ that Computation 154 reports as 2.1×10^{-24} at the horizon. *A correction carried with it*: the per-pair divergence is $-14\sigma\sigma_\nu$ and so not zero in flat space either, meaning flat conservation was always a property of the ensemble; measured here, the residual falls with the pair count at an exponent of -0.35, consistent with a finite-sample effect. *Superseded in part by Computation 163*, which supplies the exact result conditionally and settles the ambiguity; the bound here is the unconditional statement and stands under it. Re-evaluated at the derived $d_0 = 2.2$ am (Computation 201): 6×10^{-43}, 6×10^{-45}, 8×10^{-59} and 5×10^{-88}, the values Chapter 7 quotes.

Computation 163 · §7.4 · [PROOF]

The action principle $\Pi_{\mu\nu}$ does have, and what the ambiguity is.

Closes what Computation 162 left, in both halves, and corrects its closing sentence: the construction needs no action of its own. *The exact result*: §7.5 identifies $\Pi_{\mu\nu}$ as the conformal part of $T_{\mu\nu}^{(\text{fields})}$, whose conservation is Noether's second theorem for S_{matter}; splitting off the trace gives (46), $\nabla^\mu\Pi_{\mu\nu} = -\frac{1}{4}\nabla_\nu T$, exactly and on any background. Verified on a Friedmann background with a massive scalar, where the trace is not constant and both sides are nonzero, component by component on shell. *The corollary*: Computation 158 forces the trace of $\Pi_{\mu\nu}$ to vanish, so the right-hand side is identically zero and conservation is *exact*, with no $(d_0/L)^2$ residue, in the sector Π describes; the electromagnetic stress tensor is confirmed traceless for an arbitrary field strength in four dimensions as the worked instance. *The ambiguity*, backlog M4: for $H_{\mu\alpha\nu\beta} = f(g_{\mu\nu}g_{\alpha\beta} - g_{\mu\beta}g_{\alpha\nu})$ the shift $g_{\mu\nu}\Box f - \nabla_\mu\nabla_\nu f$ is symmetric, has trace $3\Box f$, and has divergence exactly $-R_{\nu\sigma}\nabla^\sigma f$, so it is identically divergence-free in flat space. *And tracelessness closes it*: $\Box f = 0$ is forced,

and a localised harmonic gauge is identically zero, by Dirichlet uniqueness on the Riemannian section (the discrete operator's smallest singular value is reported, so uniqueness is measured) and by the domain of dependence in Lorentzian signature, tested with the data nonzero outside the region so the shrinking cone is not a vacuous check. *What is left*: the conformal identification, which both halves lean on and which §7.5 already names.

Computation 164 · §7.5 · [PROOF]

The conformal identification: what the pair sum can represent, and what it cannot.

Attacks the single residue Computation 163 left, and replaces the caveat with a criterion. *Sharp in both directions*: every term of (40) is a rank-one dyad times a weight $W(s) = -|s|$ that never changes sign, times a bond kernel that is a density, so $\Pi_{\mu\nu}$ is always semidefinite; and conversely every semidefinite target is a dyad sum by its own eigendecomposition. The reachable set is *exactly* the semidefinite metric-traceless cone. Confirmed over six deliberately awkward ensembles (flattened to a plane, stretched along time, two clumps, laid on a light cone, collapsed onto a line): not one negative eigenvalue. *What is inside*: isotropic radiation and a null plane wave, hence Computation 158, Computation 159, Computation 160 and Computation 161 entire, so nothing claimed here is touched. *What is forced*: $\Lambda g_{\mu\nu}$ is indefinite, so the vacuum term's separation into its own sector is structural and not a design choice. *What is outside*: a static magnetic field, $T_{\mu\nu} = \mathrm{diag}(u, u, u, -u)$ with $u = B^2/2$, metric-traceless and indefinite, the negative eigenvalue being the tension along the field lines; an electric field likewise. A single-signed bond pushes or pulls throughout, and a magnetic field does both at once. *The gauge*: Computation 163's constant metric-traceless shift restores semidefiniteness, so what fails is the identification *as stated*, at the price of a universal constant that must dominate the largest tension anywhere and that nothing fixes, ρ being a quotient map with no empty region to pin it against. *Recorded so it*

is not repeated: a first pass argued the vacuum forbids that shift, which is wrong for exactly the surjectivity reason. *Scope, from Computation 199*: all of this is the pair sum itself; the identification is of conformal parts, and the conformal part of the sum reaches the magnetic and electric tensors and every traceless tensor obeying the null energy condition.

Computation 165 · §8.3.6 · [PROOF]

The photoelectric effect: the kinematic half, and the vertex that is still missing.

Computation 160 and Computation 161 called the effect absent outright. Half of it is not, and the halves separate cleanly. *Kinematics need the ladder and nothing else*: with the field's energy in steps of $\hbar\omega$, energy conservation gives (70), hence a threshold at $\omega_0 = W/\hbar$, a slope of exactly $\hbar$, and $\mathrm{d}K/\mathrm{d}I = 0$, all three symbolic. *The classical alternative fails on all three*: $K = I\sigma t - W$ contains no ω, so there is no threshold at all, and the energy grows with intensity. *And it fails quantitatively*: to meet Lawrence and Beams' 3×10^{-9} s bound [211] a classical field needs above $10^{10}\,\mathrm{W\,m^{-2}}$ against the $1\,\mathrm{W\,m^{-2}}$ of ordinary photoemission, ten orders. *Checked as numbers*: $\lambda_0 = hc/W$ gives 579, 525, 286 and 219 nm for caesium, sodium, zinc and platinum, each within a nanometre of tabulation. *And the picture's own end is predicted*: the Keldysh parameter reaches unity at $1.14 \times 10^{14}\,\mathrm{W\,cm^{-2}}$ for 800 nm and 13.6 eV, the measured multiphoton-to-tunnelling crossover, so single-photon absorption is an intensity regime and not a law. *The tally*: of five ingredients PST posits none. Four are outright, the ladder, the masslessness, the commutator's $\hbar$ and the electron's fermionic character. *Corrected by Computation 166*: this computation called the fifth, the vertex, not supplied. Its *form* is forced by gauge covariance (§13.3.1) and only its *strength* is scoped, so the tally is five of five with one scoped and the residue is a cross-section calculation rather than a missing ingredient.

Computation 166 · §13.3.1 · [PROOF]

The interaction vertex: forced by gauge covariance, and the third "absent" that was not.

Computation 165 called the vertex the one ingredient PST does not supply. §13.3 already derived minimal coupling from gauge covariance; this turns that phrase into algebra, with explicit Dirac matrices and a completely general spinor. *The variation is one term and it names the current*: under $\psi \mapsto e^{i\alpha(x)}\psi$ the free Dirac Lagrangian changes by exactly (156), carrying the *vector* current and nothing else, so only a counterterm on that current can cancel it. *The transformation law is an output*: requiring the cancellation returns $A_\mu \mapsto A_\mu - e^{-1}\partial_\mu\alpha$, uniquely. *The axial alternative fails twice*: it carries a different current, and (157) gives $\partial_\mu(\bar\psi\gamma^\mu\psi) = 0$ against $\partial_\mu(\bar\psi\gamma^\mu\gamma^5\psi) = 2im\bar\psi\gamma^5\psi$, both checked on the explicit spinor, while $\partial_\mu\partial_\nu F^{\mu\nu} = 0$ identically means only a conserved current can source Maxwell. *Scoped rather than absent*: the strength, through $e^{-2} \propto r_0^2$ and $Z^2 = \mathrm{e}^{-1}$ modulo the emergence premise. *And the pattern worth recording*: three times in this line of work a claim of "absent outright" was written without opening the chapter that already had it, at Computation 161, Computation 165 and here. No gate catches that.

Computation 167 · §7.5 · [PROOF]

What fixes the absolute scale of $\Pi_{\mu\nu}$, and what does not.

Computation 164 left ρ a quotient map, so realisation points exist everywhere and $\Pi_{\mu\nu}$ is nonzero in what we call empty space. The obvious next question is what vacuum density that predicts. *It predicts none*: with the weight forced (Computation 158) and the bond kernel normalised (Computation 159), the only remaining freedom in the magnitude is the pair density, measured here to carry ρ while leaving w flat across a factor of eight in the point count. *And the theory does not supply that density*: d_0 is a coherence length, which

§15.3.1 says is not the grain, and ℓ_* enters as a momentum cutoff on field modes, so reading the spacing off one rather than the other moves the implied density by $(d_0/\ell_*)^6 \sim 10^{61}$ at the bound. *Nothing standing is affected*, which the second check verifies rather than assumes: rescaling the ensemble moves ρ by orders while w, the trace ratio and the eigenvalue pattern do not move at all, because every result this sector has produced is a ratio, a shape or a count. *Two things follow*: the conformal identification carries a second gap beside Computation 164's cone, its magnitude within that cone being free; and there is *no* vacuum-energy catastrophe here, because the 10^{120} problem needs a computed scale to disagree with measurement and this construction does not produce one.

Computation 168 · §1.7 · [PROOF]

The extremal selection principle is not an extra commitment.

§1.7 called the one remaining principle "a natural choice given P1–P3 but not derived from them", and Computation 148 had already cut three to one. This closes it, and the argument is *rigidity rather than size*: choosing a proper sub-object is not a smaller commitment than taking the whole but a larger one, since one must also say *which*, and P1 supplies no selector. *The cost is measured, not asserted.* The octonions are built by Cayley–Dickson and then tested (norm-multiplicative to 4×10^{-15}, alternative to 10^{-14}, non-associative), $\mathrm{Der}(\mathbb{O})$ comes out 14-dimensional and closed under the bracket, and orbit dimensions follow as $14 - \dim(\text{stabiliser})$: a unit imaginary costs 6 ($S^6 = G_2/SU(3)$), a quaternionic triple costs 8 ($G_2/SO(4)$), the whole of V_7 costs nothing. The 6 and the 8 are textbook, so they check the arithmetic rather than feed it. *Exactly two choices are free*, the zero subspace and V_7, and P1's two clauses settle it: sole-primitive removes every intermediate choice, distinction-exists removes the trivial one. *The envelope goes the same way discretely*: of sixteen subsets of the classification only the empty and the full one need no naming datum, and the full one is the triple product at

dimension 64, recovering Computation 148. *And the criterion is not "largest" in disguise*: under $SU(3)$ a one-dimensional subspace is rigid, so rigidity is group-relative and the two agree here only because P1 leaves the symmetry unbroken. *Scope*: the Hurwitz classification and the norm-multiplicativity that puts the substrate in it are upstream and untouched.

Computation 169 · §3.4 · [PROOF]

The atom criterion, and the step that does not close it.

Computation 152 reduced §19.9.1's standing *condition* to "a coordination premise, weaker than independence". This sharpens that residue and, in an earlier version, claimed to close it; the closing step is *withdrawn* here and the refutation is published as part 3. *What stands*: balance needs $v(a) = -\sum_{b \neq a} v(b)$, one point fixed by the others, so Pr[balance] is exactly the expected conditional mass on it (planted atoms of mass p reproduce $\Pr = p$ to 0.004); hence one *atomless* conditional anywhere makes balance null, and the reach of that is measured: leaving a single property unpinned while five stay maximally coordinated turns 31% into a null set with $\Pr[|\sum v| < \epsilon] \propto \epsilon^6$, the exponent read from a cancellation-free closed form and cross-checked against Monte Carlo within five Poisson σ. *What is withdrawn*: the step pricing the atom as a distinguished direction. It priced an *unconditional* atom at a *fixed* point while the chain delivers a *conditional* atom at a random one, and Computation 152's $\mu_{\pm}$ has rotation-invariant law, names no direction, pays nothing, carries maximal conditional atoms and balances 31% of the time. *So the condition stands*, and closing it needs an argument about correlation rather than about direction.

Computation 170 · §19.9 · [PROOF]

F4's spectrum half: the alignment branch fails for the terms tested.

F4 says the theory below M_* *is* the Standard Model. Half is a matching to observation and cannot be proved; this takes the other half against the objection that a Landau–Ginzburg transition should deposit six massless modes below M_*. *It turns on a reading of P3.* With a *linear* locking term the vacuum solves exactly, giving $m_r^2 = \kappa/r + 8br^2$ against $m_t^2 = \kappa/r$, so $m_t < m_r$ at every κ with the ratio saturating at $1/3$; a *quadratic* term gives 2κ against $4(a + \kappa)$, saturating at $1/2$, and both closed forms are checked against the Hessian. Since the radial mode is the Higgs and $M_* = 10.26\, m_h$, those modes fall *below the Higgs mass. But that is not an exclusion of the branch*, which an earlier version claimed: the sign turns on the *curvature* of the locking term, and a *quartic* one, $-\kappa\langle\psi, \hat{\tau}\rangle^4$, reverses it above a crossover, with $m_t^2 = 2a\kappa/(b - \kappa)$ unbounded as $\kappa \to b$ and the angular modes *above* M_* at $\kappa \approx 0.2495$. *Reading the locking as direct determination*, which is what P3 says, leaves the order-parameter sector one mode: six generators move the vacuum and fifteen fix it, the six broken *explicitly*, so neither yields a Goldstone. *The signature theorem is insensitive*: well-posedness is decided by the principal symbol, so Computation 2 needs hyperbolicity rather than masslessness.

Computation 171 · §9.2 · [PROOF]

θ's provenance: the phase survives determination, the direction does not.

Computation 170 read P3's locking as determination without excluding the alternative, and took away the mode §9.2 was using. This says which degree of freedom θ is. *The phase*: $\hat{\tau}$ is a *real* direction, so determination fixes a direction and says nothing about a phase, and the complex line $\mathbb{C}\hat{\tau}$ is stable

under $\psi \mapsto e^{i\alpha}\psi$ where the real line is not. A real order parameter therefore has *no* candidate for θ, which makes the complexification a requirement rather than the richer of two formulations. *Masslessness is a condition, measured not assumed*: at an exactly solved vacuum the phase curvature sits at the finite-difference floor ($\sim 10^{-8}$, which is ϵ/h^2) for $|\langle\psi, \hat{\tau}\rangle|$, its square and its fourth power, and is lifted to 0.27 by $\mathrm{Re}\langle\psi, \hat{\tau}\rangle$. *Gauge invariance is what forbids the lifting term*, the moduli being invariant and the real part not, which is the same exactness Computation 170 used for the photon. *And θ is not the photon*: as a free massless scalar it would break F4 alone, and a gauged broken U(1) gives a massive boson, so θ is eaten and the photon is the unbroken combination. One flat direction, not six. *Left open*: how two substrate degrees of freedom embed in the projected four-component doublet.

Computation 172 · §8.2.1 · [PROOF]

θ is eaten by the Z, not gauge-fixed away by $\mathrm{U}(1)_{\mathrm{em}}$.

The PPN enumeration removed θ from the infrared by calling it a $\mathrm{U}(1)_{\mathrm{em}}$ gauge redundancy, Stueckelberg-eliminable because that symmetry is unbroken. *The premise is right and the inference is not*, and the same fact settles both: $Q = T_3 + Y/2$ annihilates the vacuum, which is why $\mathrm{U}(1)_{\mathrm{em}}$ is unbroken, *and* annihilates H^0, which carries charge zero, so the transformation does not move θ and has nothing to remove. *What removes it*: the phase direction $i\langle H\rangle$ lies in the span of the broken generators acting on the vacuum, of rank three, so θ is an ordinary eaten would-be Goldstone taken by the Z, which is what Computation 171 concluded from the substrate side independently. *And the list counted it twice*, once correctly among "the three electroweak eaten Goldstones" and once again as a fourth item. *The conclusion stands*: four real components, three eaten, one massive Higgs and no massless scalar, so no long-range scalar survives and $\gamma = \beta = 1$ is untouched, only the reason changing. *And the embedding closes*: the substrate's complex line is the *neutral* component, modulus to the Higgs and phase to the Z-eaten Goldstone, with

$H^{\pm}$ supplied by $\mathrm{SU}(2)_L$, which is the projection question Computation 171 left open.

Computation 173 · §12.2 · [NUMERICAL]
The seagull solved self-consistently, and b as an output.

§12.2 records that M_* and Z^2 are not jointly established, and states it with b held at $1/4$. *But if the seagull must use the canonical coupling, b is an output, not an input*, and the system closes: $m_h^2 = 6\lambda_{\mathrm{SM}}(M_*)M_*^2/(16\pi^2)$ together with the bridge $\lambda_{\mathrm{SM}}(M_*) = b\,\mathrm{e}^{-1}$ and Computation 123's two-loop running is two equations in two unknowns. *Solving it*: a self-consistent point exists at $M_* = 2203$ GeV with $\lambda_{\mathrm{SM}}(M_*) = 0.0851$, the ratio to the postulated 1285 GeV being $1/Z$ identically; and it returns $b = 0.2313 \pm 0.0020$ against P3's 0.25, the uncertainty propagated by *re-solving* for M_* at each input variation rather than holding it fixed. The gap is 7.5%, which is 9.5σ on that budget, so input uncertainty cannot absorb it. *Why that is not yet a falsification*: the coefficient C_H depends on the *shape* of the ultraviolet regulator and not only on its scale, and the book supplies a scale while never specifying a shape. Measured here: a hard cutoff and a Gaussian both give $\kappa = 1$, a smooth step of width $1/8$ gives 1.05, width $1/4$ gives 1.20, an exponential profile gives 2.00. Closing the gap needs $\kappa = 1.081$, inside that range. *What would settle it*: §12.2 already specifies the substrate spectrum (Walsh eigenmodes over non-empty $S \subseteq [D]$, masses $|S|^{1/2}M_*$), so the mode-sum envelope *is* the missing profile and is computable. It would either recover $b = 1/4$, turning P3's quartic from a postulate into a result, or falsify the bridge at a measured significance. Computation 174 carries out that calculation and returns the second outcome.

Computation 174 · §12.2 · [NUMERICAL]

The substrate regulator envelope, and what it does to C_H.

Computation 173 declines to call its 9.5σ gap a falsification because C_H depends on the ultraviolet regulator's *shape* and the book supplies only a scale; it then notes that §12.2 already specifies the substrate spectrum, so the envelope is computable and owed. This computation pays that debt. *The question to settle first* is which of the theory's two cutoffs the seagull answers to: it regulates 4-momentum at M_*, while eq. (134) regulates the substrate Dirac spectrum at the matched scaling $\Lambda = \sqrt{D}\, M_*$, a factor $\sqrt{D}$ away. *It does not have to be settled*, because both are computable and both give $\kappa = 1$. *Route A, at M_**: §12.2's mass spectrum is $m_S = |S|^{1/2} M_*$ over non-empty S, together with the zero mode, so $(0, M_*)$ is empty and the first level sits exactly at M_* with degeneracy D; a loop that runs out of modes at a spectral gap edge is cut off by a step, and $\kappa = 1$ with no D-dependence. *What that rests on*, corrected on a referee pass: not the gap alone. The Standard Model lives below M_* where the substrate has no state, so the emergent low-energy modes are collective and a substrate gap does not by itself truncate an emergent loop; the step follows once Computation 105's layer boundary is identified with the cutoff, which is architecture the book carries rather than a consequence of the Boolean Laplacian. This is the architecture Computation 105 states from the other side, the two layers occupying the disjoint ranges $[m_h^2, M_*^2]$ and $[M_*^2, D]$ and meeting at a single scale of measure zero. *Route B, at $\Lambda = \sqrt{D} M_*$*: the Walsh ladder does furnish a smooth envelope, and although the substrate's binomial density of states $\binom{D}{j}$ and the four-dimensional continuum's $\rho(j) = 2j/D^2$ look nothing alike, one peaked at $D/2$ and the other at its endpoint, the seagull sees only the first inverse moment, where $\langle 1/j \rangle_{\text{cont}} = 2/D$ exactly and $\langle 1/j \rangle_{\text{sub}} = (2/D)(1 + 1/D + O(D^{-2}))$; hence $\kappa_{\text{sub}}(D) = 1 + 1/D$, computed in exact rational arithmetic and cross-checked by log-binomial weights to $D = 65536$. *The shape that envelope corresponds to*: the family $f(x) = e^{-x^p}$ has $\kappa(p) = \Gamma(1 + 2/p)$ in closed form, and inverting it gives $p_{\text{eff}} \to 2$ with

$(2-p_{\text{eff}})D \to 2/(1-\gamma_{\text{E}}) = 4.73054$, so the substrate's envelope *is* the Gaussian that eq. (134) already uses. That choice was previously defended only by the observation that $f(1) = \mathrm{e}^{-1}$ for every p, which carries no shape information, because the substrate Dirac spectrum is degenerate and $f(1)$ is the only value evaluated; the mode sum supplies what that argument could not. *The verdict*: $\kappa = 1.081$ corresponds to $p = 1.7178$ and to κ_{sub} at $D \approx 16$, against a substrate cardinality V/a_n^4, so Computation 173's shape defence is withdrawn and the $b = 0.2313$ against 0.25 gap stands at 9.5σ. A tamper test confirms the functional is sensitive to the density of states: a uniform one over the same range gives $\kappa = H_D/2$, unbounded. *Scope*: the 9.5σ is a statement about the canonical-Standard-Model-coupling reading only; the reading in which $\lambda_{\text{PST}} = 1/4$ is canonical gives $M_* = 1285$ GeV and Computation 123's 1.0σ match and is untouched. What is excluded is holding both. Net: the open register loses an escape hatch and gains a derived regulator shape. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 175 · §14.2 · [PROOF]

Colour and generation cannot be independent labels on $\hat{\tau}^{\perp}$, and the Dixon ambient does not change that.

The colour-versus-generation caveat was closed at v25.8 and reopened at v32.64 when Computation 106's block-diagonality claim was withdrawn. The closure, and §13.5 until this computation, held that PST resolves the colour-versus-generation caveat "on the Dixon side" because the ambient $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ makes the three generation sectors orthogonal (Computation 48). *That inference is invalid, and this file shows it in three measured steps. (1) Orthogonality is not invariance.* The three Fano lines through $\hat{\tau}$ are mutually orthogonal to machine zero, are genuine projectors, and their ranks sum to $\dim \hat{\tau}^{\perp} = 6$, so

Computation 48's decomposition is real; and colour still leaks out of them at O(1). Both hold at once, because orthogonality of subspaces says nothing about whether a group preserves them, and the v25.8 closure inferred the second from the first. *(2) The obstruction is a theorem, not a dimension count.* The commutant of $\mathrm{Stab}_{G_2}(\hat{\tau})$ on $\hat{\tau}^{\perp}$ is computed to be exactly 2-dimensional over $\mathbb{R}$, that is $\mathbb{C}$, so by Schur $\hat{\tau}^{\perp}$ is the *irreducible* fundamental **3** and has no proper invariant subspace whatever. A colour-invariant generation splitting inside $\hat{\tau}^{\perp}$ is therefore impossible rather than merely unproven: the register's count rules out a tensor factorisation ($\mathbf{3} \otimes \mathbf{3}$ needs nine complex dimensions against three), and irreducibility additionally rules out the direct-sum splitting the Fano lines actually realise. *(3) The ambient changes nothing.* Colour acts on $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$ only through the $\mathbb{O}$ factor, so $\|(1 - I \otimes P)(I \otimes g)(I \otimes P)\| = \|(1 - P)gP\|$ and the leakage on the 64-real-dimensional ambient equals that on $\mathbb{O}$ alone, measured to 0 difference. The enlargement supplies room and no invariance. *What survives*: the *count* three is untouched, $\hat{\tau}^{\perp}$ being 3-dimensional over $\mathbb{C}$, and that one fact is the source of both $N_c = 3$ and $N_{\mathrm{gen}} = 3$. PST derives one three, not two, and inherits the octonionic programme's question rather than closing it. Tamper-tested: a generic element of the maximal torus, which does preserve the lines, leaks 1.3×10^{-16} against the O(1) of generic colour, so the measurement discriminates.

Computation 176 · §12.2 · [PROOF]

The field-normalisation item is the value of one number, and one of its two readings is excluded.

§12.2 carried the book's sharpest open item in a shape with no obvious line of attack: (128) presupposes a canonically normalised scalar, (132) makes Z^2 a field-strength factor, at most one of the two fields can be canonical, and the book did not say which. *The observation*: canonical normalisation is not a matter of interpretation but a statement about a kinetic coefficient, and P3 carries one explicitly, the coefficient c of the Boolean Dirichlet form,

constrained only by $c > 0$. It is not absorbable, the form being a bare edge sum with a fixed normalisation, and rescaling $\phi = \sqrt{c}\,\psi$ moves it into the quartic rather than removing it: $\lambda_{\text{canon}} = 1/(4c^2)$, hence $C_H = 3/(2c^2)$ and $M_* = 4\pi m_h c\sqrt{2/3}$, linear in c. *Everything follows from that line.* The book's $C_H = 3/2$ and $M_* = 1285$ GeV are the $c = 1$ values, the two readings are exactly $c = 1$ and $c^2 = \mathrm{e}$, and their M_* ratio is $\sqrt{\mathrm{e}} = 1/Z$, so the "$1/Z$ identically" the section asserts is *derived* rather than observed. *The bridge premise is then one equation in one unknown*, $\lambda_{\text{SM}}(4\pi m_h c\sqrt{2/3}) = \mathrm{e}^{-1}/(4c^2)$, and it is not circular: the e^{-1} is (134)'s $f(1)$, exact and D-independent with the Gaussian f derived in Computation 174; the $1/4$ is P3's; λ_{SM} runs from the measured inputs. The residual is monotone, so the root is unique, and propagating the budget by *re-solving* for M_* as Computation 173 requires gives $c = 0.9978 \pm 0.0054$, 0.4σ from the canonical value and 121σ from $\sqrt{\mathrm{e}}$. *So the second reading is withdrawn*, excluded here at 121σ, by Computation 173 at 9.5σ as a conflict with P3's $b = 1/4$, and by Computation 174 closing its only defence. *What this does not do*: it does not derive $c = 1$ from P1–P3, which is carried as a stated assumption instead, and the 0.4σ must not be read as discharging it. That agreement is Computation 123's $\lambda_{\text{SM}}(M_*) = \mathrm{e}^{-1}/4$ in another variable, one coincidence and not two, and the two readings of it carry different error budgets rather than disagreeing: Computation 123 states it over λ_{SM}, where $b\,\mathrm{e}^{-1} = 0.0920$ against the two-loop-run 0.0927 is a rigid 1.0σ, while re-solving for M_* states it over c, where it is 0.4σ; and spending it to fix c would turn a parameter-free prediction into a one-parameter fit. It is a consistency check, and a real one, the data selecting the natural value out of a continuum. *A derivation would come from* §7.8, where c multiplies the form taken to the Laplace–Beltrami form on M: either the limit fixes c, or c is absorbed into the $a_n \to 0$ scaling and is not independent. Corrected on a referee pass of this file: the first version held M_* fixed while varying c and quoted 0.9961 ± 0.0038 and 174σ, which is the shortcut Computation 173 exists to forbid. Tamper-tested: M_*/c is constant across a factor six in c while a control with $C_H = 3/(2c)$ breaks that invariance. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates

it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 177 · §7.8 · [PROOF]

Does the Mosco limit fix P3's kinetic coefficient? No, and c is not a constant.

Computation 176 reduced the field-normalisation item to P3's kinetic coefficient c and named the route that would close it: c multiplies the Boolean Dirichlet form, §7.8 carries that form to the Laplace–Beltrami form, and the target has a normalisation of its own. *The route returns a negative answer*, with corrections worth more than the negative costs. *What the limit fixes*: (21) defines the Boolean gradient as a *bare* sum over D, and (57) divides by $|D|$, so the emergent kinetic coefficient is $K = c\,|D|\,d_0^2/16$, with no free constant anywhere in the map. *What it does not fix*: canonical form in four dimensions is reached by rescaling the field, which can be done for any K, so the target imposes no condition on the value. One free number in, one out. *The consequence for c*: A1 sends $|D| \to V_M/a_n^4 \to \infty$ with d_0 fixed, so a constant c makes K diverge as a_n^{-4} and the projected field infinitely stiff. c **therefore carries a** $1/|D|$, equivalently $c \sim a_n^4 \to 0$, and the $|D|$-independent quantity P3 owns is $\tilde{c} = c\,|D|$, giving $K = \tilde{c}\,d_0^2/16$. So $c = 1$ is not an underived statement but a malformed one, and c carrying ε's dimensions makes it doubly so. *Second correction*: (57) was missing the $1/V_M$ that A2 carries, an expectation against a probability measure having been equated with an unnormalised integral; measured on an explicit four-dimensional lattice and recovered to 0.3% across a factor 81 in volume. *Third correction*: Computation 176's $\lambda_{\mathrm{canon}} = 1/(4c^2)$ is a substrate-variable identity that skips the projection, the four-dimensional coupling being bV/K^2. *Its conclusions survive intact*, because the $\sqrt{\mathrm{e}} = 1/Z$ ratio and the exclusion of reading B use only the *ratio* of two couplings and never the formula. *Corrected on a referee pass, and the correction is to this*

file: v32.74 read the postulate as leaving the gradient's normalisation unstated and added a per-site definition. (21) had already defined it, as the bare sum, and (22)'s Efron–Stein equality holds only for that form, so the addition both duplicated a definition and falsified a cited inequality. It is withdrawn; the bare reading is the book's, and c scales. *What would close the item*: a value for K, needing d_0, which Chapter 15 bounds at $\lesssim 2.3$ nm without predicting. Tamper-tested: the divergence is exhibited as a property of the bare sum and absent once the $1/|D|$ is carried, so the test discriminates rather than assuming.

Computation 178 · §3.5 · [PROOF]

The non-balance counterexample is forced rather than chosen, and it does not scale.

Computation 169 withdrew the step that would have closed non-balance, because Computation 152's $\mu_\pm$ is rotation-invariant, names no direction, pays no parameter and still balances 31% of the time at $|D| = 6$. *This computation does not close the item either*, and says so four times; what it does is measure how far that obstruction reaches. *The reduction*: on the all-parallel family $v(a) = \epsilon_a u$ the shared direction drops out, $\sum_a v(a) = (\sum_a \epsilon_a)u$, so balance is a statement about the sign vector alone. *The step that has not been used before*: §3.5 already asserts $\mu_\pm$ invariant under per-property complementation $v(a) \mapsto -v(a)$, which on this family is $\epsilon_a \mapsto -\epsilon_a$, and the single-coordinate flips act *simply transitively* on $\{\pm 1\}^{|D|}$. A law invariant under a transitive action is uniform, so the signs are forced independent and fair: $\mu_\pm$'s sign structure is the *only* complementation-invariant option rather than a modelling choice. *Hence the probability is exact*: $\binom{|D|}{|D|/2} 2^{-|D|}$, which is 5/16 at $|D| = 6$, reproducing Computation 152's measured 31% as a forced value; it behaves as $\sqrt{2/\pi|D|}$, giving some 10^{-40} at a cardinality of 10^{80}, and is *exactly zero at every odd* $|D|$, a sum of an odd number of ± 1 being odd. *And a bound with no family assumed*: for any complementation-invariant law on $(S^6)^D$,

flipping one property sends $\sum_a v(a)$ to $-2v(a) \neq 0$, pairing each balanced configuration with an unbalanced one of equal measure, so Pr[balance] $\leq 1/2$, attained at $|D| = 2$. *Scope, stated because the file is useless if this is misread*: none of it derives non-balance. A probability bound speaks about a law and not about the single configuration the substrate has, whose standing, as §3.5 makes explicit, is stability rather than probability; the vanishing is established for the family the counterexample inhabits; and 1/2 forbids certainty, not balance. The register entry stands unchanged. Tamper-tested: dropping complementation invariance while keeping rotation invariance, relabelling invariance and zero parameter cost, by taking the signs uniform on the zero-sum set, restores balance with probability *one*, so the bound is carried by that one symmetry and by nothing else, which is the assumption to challenge if anyone reopens this.

Computation 179 · §7.5 · [PROOF]

M7 restated: the realisation-point density is supplied, and it diverges.

Computation 167 ruled that the theory supplies no realisation-point density, so $\Pi_{\mu\nu}$'s magnitude within Computation 164's cone is free. *That search was over lengths*, d_0 and ℓ_*, and neither is a spacing; A1 names one outright. §18.1 introduces a_n as "the substrate site spacing, distinct from the fixed coherence length d_0", with $|D_n|a_n^4 \to V_M$, $a_n \to 0$ and local node density $1/a_n^4 \to \infty$, and (40) sums over pairs of the *same* substrate. *Corrected by Computation 182*: this row read "through the same realisation map, $\rho : D \to M$ being defined once", and the book defines *two* maps of incompatible type under that one name. No number in this row moves. *So the density is not missing; it is infinite.* With the pair cutoff the book itself uses, $|s| \lesssim d_0^2$ at the fixed coherence length (Computation 159), a cutoff that does not scale meeting a density that does gives $\Pi_{\mu\nu} \sim |D|^2 \sim a_n^{-8}$; measured here, reproducing Computation 167's N^2 and carrying it to the exponent in a_n that Computation 167 did not

take. *Which leaves a fork the book has not chosen*: if D is A1's sequence the magnitude is *divergent* rather than free and a *normalisation* is owed, measured to be $1/|D|^2$, per pair rather than per site, and therefore fixed; if D is one physical substrate then Computation 167 stands entire and a *density* is owed. Computation 167 took the second without naming it, and its arithmetic is reproduced here exactly and not disputed, its scale-free checks included. *Computation 182 decides the fork in favour of the first*, the second being unavailable rather than merely unstated, and corrects what the first buys: the per-pair normalisation converges, but the bare sum is dimensionless, so restoring the dimensions of an energy density returns one power of the node density and the object still diverges. *This does not close M7*: what is left is one truncation scale a_*, which Computation 167 showed is neither d_0 nor ℓ_*. Computation 167's negative result survives, no absolute scale being *computed* either way and so no vacuum-energy catastrophe. *The pattern is worth more than the item*: Computation 177 found the same unstated choice in P3, where a bare sum over D met a limit dividing by $|D|$, and (40) is a bare sum over *pairs* carrying no normalisation at all. The book uses D both as A1's sequence and as a fixed set, and every bare sum over D is sensitive to which; Computation 182 settles which is meant here. Tamper-tested: tying the cutoff to the spacing instead moves the measured exponent in $|D|$ from 1.98 to 1.02, so the divergence is carried by the cutoff's relation to the spacing and not by the pair sum.

Computation 180 · §10.3.2.5 · [NUMERICAL]

Born agreement: the atom obstruction survives on admissible graphs.

Computation 145 measured the Born condition's real obstruction as *atoms*, ties in T_{ex} that no monotone map can separate and that floor the distance from uniform at $m/2$, and showed the floor vanishes for generic distinction graphs. It then set the remaining graphs aside on the ground that the

book already excludes them, T_{ex} failing complement-asymmetry on regular (D, δ). *That ground does not hold, and this computation is mostly about why.* (16) is a condition on v: the body states its genuine content is the non-balance $\left|\sum_a v(a)\right| > 0$, so it constrains no distinction graph at all; and on the magnitude reading the complete graph is regular and complement-asymmetric at once, measured here. *It also names the wrong class.* The regular graphs' floor *vanishes*, as $|D|^{-0.5}$, slower than generic but vanishing. What fails is a *bounded active part*: a star has one centre, K_k with isolated vertices a fixed clique, and in both T_{ex} takes O(1) distinct values however large $|D|$ grows. *And they are admissible*: (14) and (15) are checked exhaustively on every family and hold, while (16) cannot reach them. The star keeps an atom of exactly $1/2$ at every $|D|$, half of all configurations omitting the centre and tying at zero, so the distance from uniform is floored at $1/4$ for arbitrarily large substrates. *The consequence*: Born agreement is jointly a condition on the law of T and on δ, and the theory constrains δ not at all in the relevant respect, so discharging it needs a genericity constraint on the distinction graph rather than a calculation. *Named since* (Computation 183): the constraint is the vertex cover number of (D, δ), whose being bounded is what "bounded active part" measures, and requiring only that δ tell the properties of D apart removes the mechanism. *Scope*: this is not a refutation of the crossing rule and nothing standing depends on it; Computation 145's generic result is reproduced here, not disputed, and a physical substrate is overwhelmingly more likely to be generic than to be a star. What changes is only what the register may claim. Tamper-tested: four hundred random monotone reparametrisations fail to beat the $m/2$ floor on the star, while the same search on an Erdős–Rényi graph goes far below it, so the bound is structural and the test separates the two families.

Computation 181 · §11.3 · [PROOF]

W4's named route cannot work, and the reason is one line.

§11.3 withdraws the earlier argument for microcausality as vacuous, it having used $\{c_a, c_b\} = 0$, which holds at every separation including $a = b$, so the spacelike hypothesis did no work. It then names what would establish W4: the substrate two-point structure having Fourier support in the closed causal cone. *That route is unavailable rather than unattempted.* The substrate is pre-causal and carries no dynamics, P1 supplying a Bernoulli product measure in which nothing evolves, so the two-point structure is an *equal-time* object and its Fourier transform does not depend on k^0; measured here to machine zero along every k^0 line. *The obstruction is then exact*: let F be k^0-independent and supported in $C = \{|k^0| \geq |\mathbf{k}|\}$, and fix $\mathbf{k} \neq 0$; the points with $|k^0| < |\mathbf{k}|$ lie outside C, so F vanishes there, and being k^0-independent it vanishes on the whole line, leaving F supported at $\mathbf{k} = 0$ alone. A k^0-independent two-point function has cone support only if it is trivial, so a static object cannot encode a light cone. *The numerics confirm the proof including its edge case*: at the physical end, where $a_n \to 0$ broadens the kernel in $\mathbf{k}$, some 82% of the spectral weight sits outside the cone, and the fraction falls to zero only as the kernel is squeezed onto $\mathbf{k} = 0$, which is exactly the trivial support the proof isolates. *What is not claimed*: W4 is not refuted and nothing downstream moves. Equal-time microcausality is untouched, $\{\psi, \psi^\dagger\} = \delta^{(3)}(\mathbf{x} - \mathbf{y})$ vanishing for $\mathbf{x} \neq \mathbf{y}$; the structural spin-statistics theorem never used microcausality; and W1, W2, W3 and W5 stand. *Where W4 has to come from* is the emergent dynamics, $iS(x - y)$ being built from Closure (3)'s projected Dirac operator with the Pauli–Jordan support coming from the mass shell, so the register's target moves to the matched-scaling limit. Tamper-tested: a mass-shell kernel puts nothing outside the cone while the substrate kernel puts a third of its weight there, and the control passes only because it *has* a k^0-dependence, which is precisely what the substrate lacks.

Computation 182 · §7.5 · [PROOF]

M7's fork is decided, and the missing number is named.

Computation 179 left the absolute scale of $\Pi_{\mu\nu}$ as a fork between two readings of D and said this book had not chosen. *It is not a choice.* D is finite at every stage, by A1's own $|D_n| \to \infty$, by the Bernoulli product over D, by the exact enumeration over $2^{|D|}$ configurations and by the balance probability $\binom{|D|}{|D|/2}2^{-|D|}$; a finite set has no uncountable quotient, so §4.10's $M = D/\!\sim_\rho$ is finite at every stage and is a smooth four-manifold at none, while (46) differentiates $\Pi_{\mu\nu}$ covariantly on a smooth M. So D in (40) is A1's sequence and the fixed-substrate branch is *unavailable* rather than merely unstated; Computation 167's arithmetic is untouched and its reading of the symbol falls. *And the forced normalisation does not buy a finite density*, which corrects Computation 179. Dividing by the pair count is forced and converges, but (40) is dimensionless, $W(s)$ and the unnormalised dyad carrying L^2 each against the bond kernel's L^{-4}, so the bare sum is a pure count and goes as $|D|^2$; measured here at exponent 2.021 globally and 2.078 with the bond kernel localised, so the locality costs a constant and not an exponent. Restoring the dimensions of an energy density then returns exactly one factor a_n^{-4}, the node density itself, and the restored object diverges as *one* power of it, fitted at 1.011 against controls at 0.049 and 2.049. *So one number is left*: the physical spacing a_* at which A1's limit is truncated, with $\Pi_{\mu\nu} \sim 1/a_*^4$, which is what a lattice spacing is. Computation 167's finding that neither d_0 nor ℓ_* is such a spacing stands, so M7 does *not* close; the gap is one truncation scale of a familiar kind rather than a density the theory was said not to supply, and the vacuum-energy conclusion is unaffected because a_* is named and not computed. *Two things are checked rather than assumed.* §10.5's bipartition rests on ρ being a surjective quotient, and needs none of it: a preimage and its complement partition the domain of any map, verified over 200 random non-surjective maps. And this file corrects Computation 179's premise sentence that the book "defines $\rho : D \to M$ once": it defines two

maps of incompatible type under one name, a surjective quotient onto a derived target and a non-surjective sampling of a given one. No number in Computation 179 moves.

Computation 183 · §10.3.2.5 · [PROOF]

The Born obstruction is the vertex cover number, and one line removes it.

Computation 180 left Born agreement waiting on "a genericity constraint on δ that the theory does not supply", naming the failing property informally as a *bounded active part. It is a graph invariant, and the proof is one line.* Let S meet every distinction; the remainder $D \setminus S$ carries none, so every C inside it ties at $T_{\text{ex}} = 0$, and under $\mu = \text{Bern}(\frac{1}{2})^{|D|}$ that is an atom of at least 2^{-c}, with c the vertex cover number of (D, δ). *That reproduces Computation 180's table exactly*: the star and K_4, K_6 with isolated vertices have $c = 1, 3, 5$ fixed while $|D|$ grows, and the complete, cycle and Erdős–Rényi graphs have c growing with it; nothing else in the table does any work, so "bounded active part" was the right diagnosis and this is what it measures. *And one line removes the mechanism.* Call (D, δ) *reduced* when δ alone tells the properties of D apart, no empty distinction set and no two the same; then $u \mapsto \{b : \delta(u, b)\}$ is injective from $D \setminus S$ into the subsets of S, so $|D| - c \leq 2^c$ and $c \geq \log_2(|D| - c)$, putting the guaranteed floor below 10^{-80} at $|D| \sim 10^{80}$. *The bound is tight rather than slack*: the extremal reduced graphs, a cover with one independent property per nonempty subset of it, sit within a factor 1.6 of 2^{-c} at every size and their atom falls as $|D|^{-1.09}$. *This does not close the item, and the reason is stated rather than skirted*: §2.2 says no further axioms are imposed on δ, and §4.10's appeal to the identity of indiscernibles is made for the fibres of ρ and cannot be stretched here, the substrate being the triple (D, δ, v) so that v discerns what δ does not. What changes is the residue's shape: failure by this mechanism at mass m needs at least $m(|D| - \log_2(1/m))$ properties sharing one distinction set, which at substrate scale is nearly all of them, so the register's

question becomes whether P1's primitive properties admit δ-duplicates, a decision about the postulate's wording rather than a calculation. *Settled since* (Computation 184, P1(iii)): they do not, and not by stipulation but because faithfulness (6) makes any such pair indiscernible in the whole triple. This row's sentence that §2.2 imposes no further axiom on δ records the state of the book when this computation was written and is superseded by that clause. *A second mechanism is searched for and not found*, over Erdős–Rényi graphs at three densities and by hill-climbing the atom over reduced graphs directly, which is measured over the graphs searched and not proved over all of them. Tamper-tested twice. Restoring the false twins to an extremal reduced graph holds its atom on 2^{-4} while $|D|$ goes from 19 to 59, so the cover and not the size carries the atom. And the hill-climb walks to a perfect matching, which is reduced and whose law is exactly Binomial($|D|/2, \frac{1}{4}$): its atom is the binomial mode mass, matching the enumeration to twelve digits and falling as $1.303|D|^{-1/2}$, the slowest-vanishing reduced family reachable and still vanishing.

Computation 184 · §2.2 · [PROOF]

Reducedness is a theorem of one coherence condition, and the quotient survives it.

Computation 183 left the Born condition on reducedness and declined to assume it, §2.2 imposing no further axiom on δ and v being able to discern what δ cannot. *It need not be assumed.* It follows from one coherence condition, *δ-faithfulness*: if $v(a) \neq v(b)$ then $\delta(a, b)$, that is, δ records every difference the substrate carries. *The chain is three steps, each proved and each enumerated.* Irreflexivity (5) makes every twin pair non-adjacent, $b \in N(a)$ while $b \notin N(b)$, checked over every graph to $|D| = 6$ with not one adjacent twin found; faithfulness contraposed then forces $v(a) = v(b)$ on such a pair; and the transposition $(a\, b)$ is then an automorphism of the *whole* triple (D, δ, v) rather than of δ alone, so §4.10's own principle applies one level below where the book

already applies it. Measured over 4000 triples: every equal-v twin pair is a triple-automorphism and no differing-v pair is, so the principle reaches exactly the pairs faithfulness leaves. *The dependency that had to be checked is* ρ, a spacetime point being a fibre and three live passages needing ρ non-injective. It survives exactly: (30) makes a within-fibre pair satisfy $T(\{a, b\}) = \tau_{\text{null}}$, and $\tau_{\text{null}} \in \mathbb{R}^+$ while (14) vanishes only on pairs carrying no distinction, so a within-fibre pair is *distinct* and a fibre is a *clique*, which by step one holds no twins. Over every graph on five and six properties, substrates with no shared distinction set carry clique modules of every size up to $|D|$ and independent modules of size at most one. What is removed is the opposite object, a pair with $|T| = 0 \neq \tau_{\text{null}}$ sitting at two *different* spacetime points with no fact to tell them apart. *One clause is not delivered, and running the file is how that was found*: Computation 183 defined reducedness with a second clause, no empty distinction set, and since two isolated properties are twins, at most one isolated property survives reduction while nothing removes the last. It costs nothing, Computation 183's proof resting on injectivity of $u \mapsto N(u)$, for which the empty set is a legal value one property may take and no two may share; the bound $|D| - c \leq 2^c$ is verified here directly on every collapsed substrate instead. *And what the premise does to* δ *is stated rather than left for a reader to find.* Contraposed, (6) says two properties carrying no distinction share their directional content, so δ contains the complete multipartite graph on the partition of D by v-value and the only freedom is inside a v-block; where v is injective δ is *complete*, which is admissible and whose atom Computation 180 measured vanishing as $|D|^{-0.48}$. *Computation 180's own counterexamples cease to exist*: a star's leaves are pairwise non-distinct, so (6) forces one v on all of them while they already share one distinction set, collapsing the star to two properties at $|D| = 5, 10, 20, 100$ alike, and K_k with isolated vertices to $k + 1$. *What this does not do is make the condition free*: a two-property substrate with δ empty and $v(a) \neq v(b)$ satisfies (4), (5) and everything else stated, and violates faithfulness, so P1(iii) is a premise and is carried as one in §19.9.1, with the order of discovery beside it: the clause is motivated by what makes δ the distinction relation of P1 at all, and its consequence for Born

agreement was found in this file afterwards. *Nor does it close the Born item*, the remaining half being a proof that every admissible δ satisfies the condition asymptotically, which Computation 183 measures and does not establish. Tamper-tested twice: without faithfulness, twin pairs whose v differs occur in quantity and Leibniz cannot touch them; with faithfulness but without Leibniz, faithful triples still carry twin pairs. Neither ingredient is redundant.

Computation 185 · §11.3 · [PROOF]

W4 is not an independent gap; it follows from the slice-wise CAR limit.

Computation 181 closed the route this book had named for microcausality and moved the target to the projected Dirac propagator. *That target is a corollary of the other residual listed in the same register item, and the two should be counted once.* W4 asks that $\{\psi_\alpha(x), \bar{\psi}_\beta(y)\} = iS(x-y)$ vanish at spacelike separation, and S is not a further object but the *causal propagator* $G^+ - G^-$. On a globally hyperbolic spacetime a normally hyperbolic operator has unique advanced and retarded Green's operators; a Dirac-type operator is first order but its square is normally hyperbolic, so it inherits them, with $\operatorname{supp} G^\pm(x,\cdot) \subseteq J^\pm(x)$, hence $\operatorname{supp} S(x,\cdot) \subseteq J(x)$ and CAR quantisation on a Cauchy slice gives microcausality outright [213, 214]. *Three of that theorem's four hypotheses this book already supplies*: $M = \mathbb{R} \times S^3$ is globally hyperbolic with compact Cauchy surfaces $\{t\} \times S^3$, which Computation 4's entry states while fixing the spin structure; its Cauchy problem is well posed, §4.12 deriving the signature from exactly that; and the projected field satisfies (120). *The fourth is the Cauchy-slice CAR limit*, which §19.9.1 lists beside W4 as the item's other residual, so the item has one residue and it is the slice. *The step the reduction leans on is measured rather than quoted.* The book says the Pauli–Jordan function vanishes outside the cone "because of the mass shell"; the 1+1 scalar case computed from its mode integral reproduces the exact $-\frac{1}{2}\operatorname{sgn}(t)\,\theta(t^2-x^2)\,J_0(m\sqrt{t^2-x^2})$ to five decimals inside the cone and sits

at 10^{-8} outside it. Change the dispersion and nothing else: a superluminal $c = 2$ leaves -1.1×10^{-1} outside the $c = 1$ cone and a Schrödinger dispersion 2.7×10^{-2}. *The superluminal control is the informative one*, having a cone of the wrong opening angle, so what is measured is the angle and not merely the presence of support. The spinor case needs no separate argument, S being $(i\gamma^\mu\partial_\mu + m)\Delta$ and differentiation not enlarging support; the derivatives are measured outside the cone as central differences of the converged function, five orders below their values inside. *A method note is kept because the first attempt gave the wrong answer*: differentiating under the integral removes the $1/\omega$ damping and leaves a quadrature cancellation floor at 10^{-4} that does *not* improve as the cutoff rises from $K = 100$ to 1600, while the function itself converges to 10^{-8} over the same range. *Two things are named rather than used quietly*: the Cauchy slice is not supplied here and remains the item's only residue; and the reduction needs the limit spacetime to be *Lorentzian*, which this book derives in §4.12 from strict hyperbolicity and from KO-dimension two, *not* from the Mosco theorem (§7.8), whose statement delivers E^{Riem} and the Laplace–Beltrami operator. A Dirichlet-form limit is an elliptic object and an elliptic operator carries no cone; reading cone support out of the Mosco limit directly would be the same elliptic-for-hyperbolic confusion Computation 181 found one level down.

Computation 186 · §11.3 · [SURROGATE]

The Cauchy slice is a slab, and the missing power of mass is its thickness.

Computation 185 reduced this item to one residue, the slice-wise CAR limit, and did not supply it. *The construction it was missing is forced by a count this book had already made.* A Cauchy slice is three-dimensional inside the four-dimensional M while $\rho_n(D_n)$ is finite, so $\rho_n^{-1}(\Sigma)$ is *empty* at every size, measured here; posing the relations slice-wise therefore cannot mean restricting to sites realised on Σ, because there are none. *The book's own diagnosis of the*

earlier error says what it does mean. That version "wrote $\delta^{(4)}(x-y)$, summing over sites realised into the four-dimensional M rather than into a slice", which is "dimensionally wrong by one power of mass", and *the missing power is a thickness*: the slab $S_n(t) = \{a : |\tau(\rho_n(a)) - t| \le h_n/2\}$ has site density $\nu_n = \lambda_n h_n = h_n/a_n^4$ per unit 3-volume, carrying $[\text{mass}]^3$ where λ_n carries $[\text{mass}]^4$, and smearing with $\nu_n^{-1/2}$ lands on $\int_\Sigma \langle f, g\rangle \, d^3x$, the distributional $\delta^{(3)}$. Measured against the exact slice integral, with the old reading beside it low by exactly $1/h$. *The error then splits in two and each condition kills one term*: $h_n \to 0$ kills a bias measured as $\mathrm{O}(h^2)$, tracking its predicted coefficient to a constant of about 1.2 on the rows where it clears the standard error of the mean, and $\lambda_n h_n \to \infty$ kills a fluctuation measured as $\mathrm{O}(\nu_n^{-1/2})$, which is the CLT rate and so the Berry–Esseen rate this section says the convergence should match. Thinning the slab helps one and hurts the other, so these are two conditions rather than one stated twice, and both hold on the wide window $a_n^4 \ll h_n \ll 1$. *Balancing them picks out a preferred thickness*, $h_n \sim a_n^{4/5}$ with total error $a_n^{8/5}$; the optimal slab is *thicker* than the site spacing and holds $a_n^{-1/5}$ sites across it. The measured RMS minimum sits within a constant of order one of that scaling and the convergence exponent at -0.36 against -0.40, both reported rather than tuned. *Tamper-tested three ways*, and the informative one is the slab that never shrinks: it converges perfectly well, to the slab average instead of the slice value, so the measurement distinguishes a slice from a slab rather than detecting noise. *What this does not do*: it is the i.i.d. surrogate, the standing of A2's verification and A3's corroboration (Computations 70, 83), and for the emergence-determined ρ_n the equidistribution it leans on *is* A3, supplied by P3 rather than here, with P3's time function needed besides. So the item is not closed here; it has gone from an absence to a construction resting on the condition the Mosco chain already carries. Stands in for: the emergence-determined ρ_n, by an equidistributed cloud in a unit 4-box.

Computation 187 · §4.1 · [PROOF]

Only the atom (14) **forces survives a generic T, and that is the one bounded.**

Every atom this book has measured in the Born obstruction, in Computation 145, Computation 180 and Computation 183 alike, was measured on T_{ex}, which §4.1 introduces as "a minimal example for the *magnitude* structure" and names as "one representative among" a family whose dimension grows as $2^{|D|}$. T_{ex} *is integer-valued, so ties are forced on it by counting alone*: at $|D| = 14$ it maps 16384 configurations onto a few dozen values. P2 makes the theory's T V_7-valued and hence $|T(C)|$ real, so asking which atoms are the theory's and which the exemplar's needs no new premise; it is the book's own freedom, used. *The split is exact.* (14) puts $|T(C)| = 0$ on exactly the configurations with no distinct pairs, which is admissibility rather than choice, so *every* admissible T carries an atom of mass μ(independent sets) and no other tie is forced: largest atom $\geq \mu$(independent), with equality for a generic one. Measured with positive real edge weights, admissible and checked here against (14) and (15) rather than asserted, the largest atom equals μ(independent) *exactly* in every family Computation 180 tabulated, over twenty independent weightings each. *A correction to Computation 180's rate, not to its verdict*: the complete graph's atom is not 0.21 falling as $|D|^{-0.5}$ but 9×10^{-4} falling exponentially, its independent sets being only the empty set and the singletons, so the passing families pass for a better reason than was recorded and §10.3.2.5's sentence is corrected. *An upgrade to Computation 183*: its proof runs through the configurations inside a cover's complement, which are independent sets, so the chain largest atom $\geq \mu$(independent) $\geq 2^{-c}$ is T-independent at both steps and holds for *every* admissible T rather than for the exemplar it was measured on. That also *explains* rather than re-measures why Computation 183's search for a second mechanism found none: under a generic T there is none, the cover bound being the whole obstruction rather than one mechanism among several. *What it does not do is remove P1(iii)*: the surviving atom is precisely the one

reducedness kills, and the star holds $1/2 + 2^{-|D|}$ under every admissible T, matched here to its closed form. Tamper-tested on the step that carries the weight: letting a third of the weights vanish breaks (14), and the atom then leaves μ(independent), so the equality is a consequence of *admissibility* and not of merely using real numbers.

Computation 188 · §9.1 · [NUMERICAL]

Neither an area law nor the finiteness fixes Λ's magnitude.

Closes a question the register carried since July 2026: whether the area-law property of Π or the finiteness of the internal structure constrains the static valley depth $\Delta\mathcal{F}$. (i) The second edition derives no area law (Table 2), so the first leg is moot. (ii) An area law quantised in a length the theory owns would miss $A/4\ell_P^2$ by $(\ell_P/\ell)^2$: 10^{-32} for ℓ_*, 5×10^{-53} for d_0 at the 2.3 nm bound of its day (5×10^{-35} at the derived 2.2 am). (iii) The energy bound such a law implies, no more vacuum energy in a region than a black hole of its size, is $3c^4/(8\pi G_\ell L^2)$ with $G_\ell = \ell^2 c^3/\hbar$; at the Hubble scale it is 10^{-32} and 5×10^{-53} of the critical density, both violated by the observed 0.7, so such a law is false, not weak. (iv) With the emergent G of (61) the bound is general relativity's, $\rho_\Lambda \leq \rho_{\rm crit}$, satisfied at 0.71 and an inequality on $j_\tau\,\Delta\mathcal{F}$, fixing neither factor; Computation 82's $j_\tau \approx 8 \times 10^{-38}$ per configuration is reproduced. (v) The finiteness sets M_*^4, $10^{59}\,\rho_\Lambda$, Mechanism A of §9.1. Negative on both legs, no fourth postulate, nothing imported; the sign, $w = -1$ and non-dilution are untouched.

Computation 189 · §8.3.6 · [PROOF]

The photoelectric cross-section and rate from the derived vertex.

The application Computation 166 left. (1) In the non-relativistic limit the vector vertex couples $\mathbf{A}$ to the electron's momentum; Fermi's golden rule on a

hydrogenic $1s$ state with a plane-wave final state gives, by a symbolic angular integral, $\sigma_B = (2^8\pi/3)\,\alpha\,a_0^2\,k^3/(1+k^2a_0^2)^5$, whose high-energy form is the Born law $(2^8\pi/3)\,\alpha\,a_0^2\,(I_{\mathrm{H}}/\hbar\omega)^{7/2}$. (2) The exact Coulomb-wave dipole result [215] tends to it as $k \to \infty$ (ratio 0.97 at $ka_0 = 100$) and supplies the threshold, 6.30 Mb for hydrogen, where the plane-wave form vanishes as k^3. (3) The axial alternative reduces to $\boldsymbol{\sigma}\cdot\mathbf{A}$, computed from the explicit Dirac matrices, a pure spin operator with zero matrix element between orthogonal eigenstates: no photoeffect at this order, so the forced form is also the needed one. (4) The rate is σ times the photon flux: at $1\,\mathrm{W\,m^{-2}}$ and 20 eV, 2.2 Mb and $7\times10^{-5}\,\mathrm{s^{-1}}$ per atom with an exponential waiting time and no accumulation, so 10^{15} atoms emit their first electron within 1.4×10^{-11} s, inside the Lawrence–Beams 3×10^{-9} s, where a classical field accumulating a work function through πa_0^2 needs 248 s. (5) α is taken as measured (§13.3.1); the law's form and its proportionality to α are the derived content.

Computation 190 · §10.3.2.5 · [PROOF]

Born agreement on every reduced distinction graph: the forced atom vanishes, as a theorem with a rate.

The theorem Computation 183 measured and the register carried as missing. Under $\mu = \mathrm{Bern}(\frac{1}{2})^{|D|}$ the forced atom is $a_0 = \mu(C \text{ independent})$, by (14) the same for every admissible T (Computation 187). (1) *Lemma 1*: a matching of m disjoint distinctions gives m independent events of probability $\frac{1}{4}$, none of which an independent C realises, so $a_0 \le (3/4)^m$. (2) *Lemma 2*: for a maximal matching the unmatched properties are independent, so each has its distinction set inside the $2m$ matched ones, and reducedness (Computation 184) makes those sets distinct: $|D| - 2m \le 4^m$. Hence $a_0 \le \max\big((3/4)^{|D|/4}, (|D|/2)^{-0.21}\big)$ for every reduced (D,δ), 10^{-17} at $|D| = 10^{80}$. (3) Both lemmas are checked exactly, a_0 enumerated over all $2^{|D|}$ configurations, on 200 random graphs at $|D| = 12$ to 16, on the subset-graph family to $|D| = 520$ by a closed form for its independent-set count, exhaustively over every reduced graph at $|D| \le 6$

and by hill-climbing at $|D| \le 16$; the largest reduced atom found sits well under the bound. Lemma 2 fails on the star and on K_k with enough isolated properties, the non-reduced families of Computation 180, so the hypothesis does the work. (4) Tamper: false twins on the $c = 3$ subset graph freeze the matching at $m = 3$ while $|D|$ grows to 73, break Lemma 2, and pin the atom on 2^{-c} to 10^{-13}. (5) The bound is universal and not sharp: the subset-graph atom falls as $|D|^{-1}$, its $|D|^{-1.1}$ over $c = 3$ to 7 reproducing Computation 183. What remains is a condition on the law of T alone, item 1's residual.

Computation 191 · §11.3 · [PROOF]

The Cauchy slab on a deterministic cloud: the slab-wise CAR limit needs equidistribution with a rate, not independence.

Removes the surrogate from Computation 186. (1) For any N points in the unit 4-box the slab estimator $S_h = (Nh)^{-1} \sum_{\text{slab}} \langle f, g \rangle$ obeys the identity $S_h - \int_\Sigma \langle f, g \rangle = \text{bias}(h) + \text{disc}(P, h)/h$: the bias is cloud-independent and exact ($3h^2/4 + h^4/20$ on the test integrand), and the second term is the cloud's equidistribution defect tested on the integrand times the slab indicator, which Koksma–Hlawka bounds by $C\, D_N/h$. So $|S_h - \int_\Sigma| \le c\, h^2 + C\, D_N/h$, the limit holds for every sequence of clouds with $D_N/h_n \to 0$, and the balanced thickness is $h_n \sim D_N^{1/3}$ with error $D_N^{2/3}$. (2) On the Halton and Kronecker sequences, eight disjoint blocks each, the error falls below the i.i.d. surrogate's at every size from 2^{12} to 2^{20} and at the predicted rate, $|D_n|^{-0.65}$ and $|D_n|^{-0.48}$ against $-2/3$, where the surrogate gives -0.39; a jittered grid, i.i.d. below its cell scale, tracks the surrogate, so stratification at one scale is not equidistribution at every scale. (3) The discrepancy term alone falls as $|D_n|^{-0.92}$ and $|D_n|^{-0.86}$ on the two sequences against the surrogate's $|D_n|^{-0.60}$, the CLT's. (4) A smooth window in place of the sharp slab tests the cloud on a C^1 function, A3's class, with bias $9h^2/7$ exactly and the same convergence. (5) Tamper: a Halton cloud with the same count but a spatial void refilled elsewhere converges to the wrong slice value, the defect flat across a sixteenfold count increase and

equal to the analytic void bias, so A1 is not the hypothesis. (6) Tamper: a cloud equidistributed in space with time density $1 + \epsilon \cos 2\pi t$ returns $(1 - \epsilon)$ times the slice value, so the time function must be the one in which the cloud is 4-volume equidistributed. What the item rests on is A3's rate and that time function, nothing of its own.

COMPUTATION **192** · §11.3 · [PROOF]

The time function the slab is cut in: isotropy forces the foliation, its level sets are the Cauchy surfaces, and the slab's normalisation is the proper 3-volume.

Supplies what Computation 191 left: the time function, built from what the book already has. P3 fixes $\hat{\tau}$ globally; Chapter 5 maps it to $\Phi(\hat{\tau}) = \partial_t$ along the $\mathbb{R}$ factor; the Cosmological Principle makes the geometry Friedmann–Robertson–Walker; A3 equidistributes the cloud against $a(t)^3\, dt\, d\Omega_3$. (1) *The foliation is forced.* A timelike field is tangent to a foliation iff $u \wedge du = 0$, iff its twist vanishes; the twist is a spatial vector, and isotropy admits no invariant spatial vector, so it is zero. Symbolically on (220), $u = dt$ has $u \wedge du = 0$ for every $a(t)$, while the rotating tamper $u = dt + \Omega r^2 d\phi$ gives $2\Omega r\, dr \wedge d\phi \wedge dt$: no foliation, no slab. (2) *The level sets are compact Cauchy surfaces*: $dt/d\lambda > 0$ along every future-directed causal curve, checked on two hundred null geodesics on $a = e^{Ht}$ and on the matter-plus-instantiation solution, each crossing each of six level sets exactly once. (3) *The normalisation is the proper 3-volume.* A3 makes the cloud's time marginal proportional to a^3, so the uniform normalisation of Computation 191 would return $a(t_0)^3/\langle a^3 \rangle$ times the slice value, from 0.2 to 2.3 across the range tested, while $\nu = \lambda\, a(t_0)^3 V_3\, h$ lands on it at every slice to 0.1%. (4) *The one choice left carries no content*: relabelling the leaves by a monotone f moves the uniform estimator by $1/f'(t_0)$ and the covariant one not at all. With Computation 185 and Computation 191, the slab-wise CAR limit, and W4 with it, is a theorem given A1 to A4 with A3 supplied by P3: the standing of equicoercivity A7, and the item leaves the

structural front on those terms.

Computation 193 · §19.9.1 · [PROOF]

One condition, two faces: the Born condition on T's contingent content and non-balance are both the non-degeneracy of that content.

The contingent content of an admissible T is a point (t, v): a magnitude t, any order-preserving function on 2^D vanishing exactly on the independent sets (a cone of dimension $2^{|D|} - i(G)$), and the directions $v \in (S^6)^{|D|}$. (1) *Magnitude face.* The law of $|T|$ under μ has an atom off zero iff two non-independent configurations tie; the tie set is a finite union of $\binom{2^{|D|}-i(G)}{2}$ hyperplanes, closed of codimension one. On generic points of the cone (increments continuous, two hundred draws on three graphs at $|D| = 6, 8, 10$) the largest atom is $i(G)/2^{|D|}$ exactly and there is no unforced atom; T_{ex} has unforced atoms, taking at most $|E| + 1$ values. (2) *Direction face.* The balance set is closed of codimension seven: for v uniform on $(S^6)^8$, $\Pr(|\sum v| < r)$ scales as $r^{6.7}$ over $r \in [0.45, 1]$ on four million draws, and $\mu_\pm$ balances with $\binom{|D|}{|D|/2}2^{-|D|}$, zero at odd $|D|$. (3) *Independence*: in the table (generic or T_{ex}) $\times$ (continuous or balanced v), unforced atoms follow the first factor only and balance the second only. (4) The distance to degeneracy, $\min\big(\min_{C \neq C'} |t_C - t_{C'}|/\sqrt{2},\ |\sum v|/\sqrt{|D|}\big)$, is positive on every generic draw and zero on every degenerate one. So both faces are the same kind of statement, open and dense, null under continuous content, positive under discrete content, fixed by the theory on neither; the register carries one condition with two faces, the magnitude face being item 1's residual, and the count is unchanged.

Computation 194 · §9.5 · [PROOF]

The unit of orbital angular momentum: $L_z = nQ$ exactly, the quantum of Q is an action the classical theory does not have, and the only action scale is the normalisation of the gradient term.

(1) From (81), $xT_{y0} - yT_{x0} = \partial_\phi\theta \cdot \Pi$ with $\Pi = 2r_0^2\,\partial_t\theta$, and for every axisymmetric $\theta = n\phi + \chi(\rho, z, t)$, $\partial_\phi\theta = n$ identically, so $L_z = nQ$ with $Q = \int \Pi\, d^3x$ the Noether charge: an identity, no far-field assumption, checked symbolically. (2) Q has the dimension of action (checked on (M, L, T) with r_0^2 an energy per length as (90) has it); the canonical commutator makes $Q \in \hbar\mathbb{Z}$ and the Bohr–Sommerfeld cycle makes $2\pi Q \in h\mathbb{Z}$, each with the unit as input; classically Q is continuous. (3) $r_0^2 d_0^2/c$ is the unique monomial in (r_0^2, d_0, c) with the dimension of action (exponents searched in $[-3, 3]$), so a derived unit $\hbar = \kappa\, r_0^2 d_0^2/c$ is $r_0^2 = \hbar c/(\kappa d_0^2)$, the absolute Goldstone stiffness, which is what the normalisation of P3's gradient term sets through $G = 6\pi d_0^2/(c\, r_0^2)$ (74): item 4's residual in another variable. Illustration, not prediction, and corrected by Computation 197: $m_0c^2 = (4\pi^3/3\kappa)\,\hbar c/d_0 = 3.5$ keV$/\kappa$ at $d_0 = 2.3$ nm was printed with r_0 left free, while the book fixes $r_0 = v$, which makes $\kappa = 1.2 \times 10^{-19}$ there and the mass 2.9×10^{13} GeV. (4) On the regulated vortex $\theta = n(\phi - \omega t)$, $L_z = nQ$ at every outer radius while both grow as R^3 and the energy per unit length as $\ln(R/d_0)$: the global vortex, a property of the class and not a gap; $\pi_2(S^1) = 0$ keeps the lump of (92) a mass scale and not a vortex. The item has nothing of its own and folds into item 4, as that item's closing clause already says.

Computation 195 · §19.9.1 · [PROOF]

Maximum entropy supplies non-degeneracy, if P3's typicality reaches the content of T.

The argument Computation 193 named as the one that would close the

condition, and what it needs. The Gibbs variational principle: relative to the invariant reference measure (the uniform product on $(S^6)^{|D|}$, Lebesgue measure on the admissible cone with a scale constraint), $D(P \,\|\, \nu) \geq 0$ with equality only at ν, and the maximum-entropy law under moment constraints is an exponential family with a density; so every closed nowhere-dense set is null under maximum entropy, the tie set and the balance set among them. (1) *Directions*: the von Mises–Fisher family of concentration κ about $\pm u$, u uniform and the signs fair (the signs Computation 178 forces), lies inside the symmetry package; its entropy, exact from the Bessel functions, decreases strictly in κ from $\log(16\pi^3/15)$ at the uniform product; exact balance has probability zero at every finite κ (a power law of slope six to seven at small radius on a million draws) and the binomial mass $\binom{|D|}{|D|/2}2^{-|D|}$ only at $\kappa = \infty$, which is $\mu_{\pm}$, of infinite relative entropy. (2) *Magnitudes*: exponential increments along inclusion, the maximum-entropy increments under a mean constraint, carry only the forced atom on 200 draws per graph; a lattice of L increment levels ties, with mass 0.27 at $L = 1$ (T_{ex}) falling to 0.01 at $L = 256$. (3) *The reference measure is load-bearing*: on the fourteen octonion units, Computation 152's quantised scope, the maximum-entropy law is uniform on a finite set and balances with probability $1/14$ at $|D| = 2$ and positively at every even $|D|$. (4) Both faces hold on every one of 500 joint maximum-entropy draws. What the argument needs: that P3's clause, which selects the realised configuration consistent with the tension data, reach the content of T that is its input. A reading of P3, not a new postulate; under it the condition has A3's standing, and the register carries it as what that reading would discharge.

Computation 196 · §7.5 · [PROOF]

The truncation spacing a_* is comoving, and under the conformal identification it is the thermal wavelength of the state's radiation.

(1) On the Friedmann geometry with the curvature k symbolic, the divergence of a homogeneous isotropic traceless tensor is $\dot{\rho} + 4H\rho$ exactly, the spatial

components zero; $\rho = Ca^{-4}$ solves it and $\rho = \text{const}$, which is a fixed physical a_*, leaves $4H\rho$; dust gives $3H\rho$, so the fourth power is the tracelessness; a perfect fluid is traceless iff $w = 1/3$. (2) Under $\Pi^0{}_0 = \rho_r$: $\rho_\gamma(2.7255\ \text{K}) = 4.17\times10^{-14}\ \text{J m}^{-3}$, neutrino multiplier 1.69, $a_* = (\hbar c/\rho_r)^{1/4} = 0.818$ mm, equal to $(15/\pi^2 f)^{1/4} = 0.974$ times $\hbar c/kT_{\text{CMB}}$ and scaling as $1/T$; $a_* = d_0$ at $kT = 83$ eV and $a_* = 0.41\,\ell_*$ at $kT = M_*$. (3) Tamper: a fixed a_* fails conservation at the relative rate $4H_0$; a general w gives $3(1+w)H\rho$. The magnitude of $\Pi_{\mu\nu}$ is supplied by the state and not owed by the theory, which was the first structural item's only remaining gap. *Scope, from Computation 201*: the 83 eV crossing is at the Casimir bound; at the derived d_0 it is $kT \approx 40$ GeV.

Computation 197 · §19.9.1 · [PROOF]

The normalisation item with $r_0 = v$: the derived-unit coefficient is 1.2×10^{-19} at the Casimir bound, and the dependency graph of K has one free node, d_0.

(1) $\hbar c/(r_0^2 d_0^2)$ is dimensionless on (M, L, T) and $6\pi d_0^2/(Gv^2)$ carries $[M]^{-2}$, as the note after (74) says. (2) With $r_0 = v$: $\kappa = \hbar c/(v^2 d_0^2) = 1.21 \times 10^{-19}$ at $d_0 = 2.3$ nm; $\kappa = 1$ at $d_0 = \hbar c/v = 0.80$ am, 2.9×10^9 below the bound; (92) gives $m_0 = 2.9 \times 10^{13}$ GeV at the bound and $(4\pi^3/3)v = 10.2$ TeV at $\kappa = 1$, linear in d_0 (1 TeV at $d_0 = 7.9 \times 10^{-20}$ m); Computation 194's 3.55 keV/κ divided by that κ is the same 2.9×10^{13} GeV, so the illustration is corrected and the arithmetic is not. (3) The graph: v, m_h, G measured; $\varepsilon = m_h^2/4$ and $b = m_h^2/(8v^2) = 0.032$ fixed by them; $c = 6\pi d_0^2/(Gv^2) = 6.3 \times 10^{48}\ \text{GeV}^{-2}$ at the bound; K, c, κ and m_0 hang on d_0 and on nothing else. (4) Tamper: with r_0 free the coefficient is unconstrained, which is the reading Computation 194 took. The item stays open as an assumption; what closes it is a prediction of d_0. *Scope, from Computation 201*: the graph omits (205), which ties K to d_0; with it the canonical reading fixes $d_0 = 2.2$ am and $\kappa = 0.13$, order one.

Computation 198 · §12.2 · [NUMERICAL]

Stability of the regulated vacuum at two loops: the running quartic is positive on the whole range the theory contains.

Computation 123's two-loop running copied and completed with the standard two-loop β_λ (GUT-normalised g_1). (1) Validation: $\lambda(M_*) = 0.09145$ at one loop and 0.09234 without the λ piece, Computation 123's numbers to five decimals; the two-loop λ terms move it by -0.00006 to 0.09228, within the 0.0007 of the literature's 0.0927; the quartic crosses zero at $10^{9.7}$ GeV, inside the known window; 0.4σ from $\mathrm{e}^{-1}/4$. (2) On $[m_t, M_*]$, λ decreases monotonically from 0.126 to 0.092 with its minimum at M_*, so $V = \lambda(\phi)\phi^4/4$ is positive at every field value the regulated theory contains; the instability scale is 6.5 decades above M_*; $\lambda(M_*) = 0$ would need $m_t = 224$ GeV; over the 3σ box of (m_t, α_s, m_h) it stays in $[0.088, 0.096]$. (3) The mass parameter's running is multiplicative ($m^2(M_*)/m^2(m_t) = 1.06$), so the sign driving symmetry breaking is the same at every scale. (4) Tamper: without $-6y_t^4$ the quartic never crosses zero below 10^{19} GeV, so the check is sensitive to the terms it carries. Stability beyond leading loop order is established in the RG-improved sense; the item is the construction at all orders. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 199 · §7.5 · [PROOF]

The conformal part of the pair sum reaches exactly the traceless tensors obeying the null energy condition, so the identification has no obstruction.

Removes the obstruction Computation 164 found. (1) Term by term, $\Delta_\mu\Delta_\nu -$

$\frac{1}{4}s\,g_{\mu\nu}$ is traceless identically (symbolic); the conformal part of the sum is the sum of these; the trace of the raw sum is $\sum |s|\, s$, minus Computation 158's summed $T - \tau_{\text{null}}$, the balance scalar. A timelike unit bond gives $\rho = 3/4$, $p = 1/4$; a spacelike one gives pressure $3/4$ along itself and tension $1/4$ across. (2) Explicit clouds: 96 points on concentric regular polygons in the x-y plane give $\text{diag}(u, u, u, -u)$ after the projection with deviation 3×10^{-16} (magnetic, field along z), the y-z plane gives the electric tensor along x; one x-bond plus one y-bond is $\frac{1}{2}\text{diag}(1, 1, 1, -1)$ exactly. In Computation 158's balanced regime the trace fraction runs 2.6×10^{-2}, 2.9×10^{-3}, 2.4×10^{-4}, 7.4×10^{-5} as the window doubles, and the projection changes the sum by a quarter of that, so Computations 158 to 161 are untouched; Computation 164's six shapes all obey the null energy condition after projection. (3) Inclusion, symbolic: for null $k = (1, \mathbf{n})$, $G(k, k) = (\Delta \cdot k)^2 \geq 0$. (4) Equality, numerical: random positive semidefinite metric-traceless Y decompose into 200 sampled null dyads with zero residual (the dual cone); random traceless tensors obeying the condition with margin decompose into 1600 sampled generators with zero residual, violators keep a residual above 0.5 times their violation, a null dyad being the certificate; the magnetic and electric tensors are reached at 10^{-3} and 7×10^{-5} while the zero-energy anisotropic stress and negative-energy radiation are not. (5) Pataki: a spectrahedron with two equations has rank-one extreme points only, so the dual's extreme rays are null dyads; the rank-two case by hand, 200 draws each equal to $c_1 k_1 k_1 + c_2 k_2 k_2$ on the two null directions of its Lorentzian plane, residual 10^{-13}, coefficients positive. (6) Tamper: a signed weight $W = -s$ reaches the zero-energy anisotropic stress (residual 0 against 1) and gives a cloud strung along the time axis a negative energy density. The identification is of conformal parts; an unbalanced sum carries a trace it does not use.

Computation 200 · §12.2 · [PROOF]

The interacting vacuum of the regulated scalar sector exists non-perturbatively; what remains is the gauge and fermion sector, the cutoff Standard Model's own.

Three classical theorems applied to a theory that never removes its cutoff. *Finite volume*: with finitely many modes the Hamiltonian is a Schrödinger operator on $L^2(\mathbb{R}^N)$ with a quartic potential bounded below, essentially self-adjoint with compact resolvent, and its semigroup is positivity improving, so the ground state is unique and strictly positive. *Infinite volume*: the Euclidean lattice theory at spacing ℓ_* has ferromagnetic couplings and an even single-site weight, so with Dirichlet data Griffiths' second inequality makes every correlation non-decreasing in the volume, and the quartic bounds it; the limit exists in any dimension, the dimension entering only the continuum limit the theory does not take. *Back to a Hamiltonian*: the nearest-neighbour action is reflection positive, so the transfer matrix is a positive operator and $-\ell_*^{-1}\log T$ is bounded below with the Euclidean vacuum as its ground state. Checks: (1) one and two sites by finite differences, every ground state nodeless and even, the double well's gap closing as it deepens (0.35 to 3×10^{-5}) and as a site is added (0.0097 to 5×10^{-5}); (2) the one-site Trotter transfer matrix positive to round-off with its gap converging to the Hamiltonian's, 0.3344, 0.3427, 0.3451, 0.3457 against 0.3458 as the step halves; (3) four-dimensional lattice Metropolis with Dirichlet data at the book's couplings in units of ℓ_*, $\lambda = 0.092$ and $\mu^2 = -(m_h\ell_*)^2/2$, the central nearest-neighbour correlation 0.0264, 0.0271, 0.0295 at $L = 4, 6, 8$ and $\langle\phi^2\rangle$ bounded at 0.15, and at a massive test point flat within errors; (4) tampers: an antiferromagnetic coupling leaves Griffiths' ordering and gives negative correlations, and a negative quartic has no ground state, the lowest eigenvalue running off as the grid widens, so $\lambda > 0$ (Computation 198) is what the construction needs. Not closed and said so: the gauge and fermion sector below M_*, whose infinite-volume limit at weak coupling and whose chiral fermions on a regulator are the cutoff Standard

Model's own open problems, inherited with it.

Computation 201 · §15 · [PROOF]

The gradient normalisation and the coherence length are one unknown: under the canonical reading, $d_0 = \sqrt{2}\,\hbar c/m_h$, and the nanometre expectation does not follow from the equations.

(205) with $|a_{LG}| = m_h^2/4$ leaves one free quantity, c_{LG}, which is the open item. (1) Symbolic: $V''(r_0) = 4\varepsilon$ for $V = -\varepsilon r^2 + br^4$ at $r_0^2 = \varepsilon/(2b)$, independent of b; the Ginzburg–Landau length $\sqrt{c/(2|a|)}$ is $\sqrt{2}$ times the radial correlation length $\sqrt{c/V''}$; at $c_{LG} = 1$, $d_0 m_h = \sqrt{2}$ in the neutral frame and 1 in the doublet frame. (2) Numbers: $d_0 = 2.23$ am (neutral), 1.58 am (doublet); $d_0/\ell_* = 14.5$ and $10.26 = M_*/m_h$; $d_0/\ell_P = 10^{17}$; the bound exceeds d_0 by 10^9; the Casimir correction $\xi(d_0/d)^2$ is 1.1×10^{-19}, 1.8×10^{-20}, 1.8×10^{-21} at 20, 50, 160 nm against 0.12, 0.019, 0.0019 at the bound and a precision of 1.9×10^{-3}; Computation 197's $\kappa = m_h^2/(2v^2) = 0.129$; (92) gives 28 TeV; Computation 196's a_* passes through d_0 at $kT = 36$ GeV. (3) A nanometre d_0 needs $c_{LG} = 1.1 \times 10^{18}$, which through Computation 176's linearity moves M_* by 10^9; even $c^2 = \mathrm{e}$ is excluded at 121σ. (4) The BCS analogy quantified, $d_0/\ell_* = \sqrt{2}\,M_*/m_h$ with $M_*/m_h = 4\pi\sqrt{2/3}$ the book's own; tamper, $|a_{LG}|$ left free: 2.3 nm needs $|a_{LG}|$ at 10^{-18} of its value, a radial mode of 10^{-7} GeV. The 0.4σ agreement is not spent: reading A is the standing premise and d_0 follows from it. *Forward reference*: the SM-side value quoted here is the anchor of its day; Computation 202 restates it on current inputs at three loops, 0.0930 ± 0.0008 on the 2026 Review, 1.4σ from $\mathrm{e}^{-1}/4$, which is what the book now states (earlier printings gave 1.6σ on the 2024 Review).

Computation 202 · §19.7 · [NUMERICAL]

The Higgs self-coupling test at three loops on current inputs: the third loop is worth nothing at M_*, the 2013 anchor is not reproducible, and the honest agreement is 1.6σ on the 2024 Review and 1.4σ on the 2026 one.

The complete three-loop Standard Model β functions and the next-to-next-to-leading-order matching conditions at the top mass, both from [51], coded and validated by reproducing that paper's Planck-scale couplings to four decimals. (1) The third loop moves $\lambda_{\rm SM}(M_*)$ by $+0.00007$, a tenth of the error, so the running was converged at two loops. (2) The anchor Computation 123 used, 0.0927 ± 0.0007, is not reproduced: the source's formulas on the source's 2013 inputs give 0.0924 ± 0.0011 at every loop order, a 0.4σ agreement, and the ± 0.0007 counted m_t at ± 0.5 GeV where 2013 had ± 0.76. (3) On PDG 2024 inputs [216] ($m_t = 172.57 \pm 0.29$, $m_h = 125.20 \pm 0.11$, $\alpha_s = 0.1180 \pm 0.0009$, $M_W = 80.3692 \pm 0.0133$) the value is 0.0932 ± 0.0006 from the quoted errors, 2.2σ, and 0.0932 ± 0.0008 once the 0.5 GeV ambiguity between the measured and the pole top mass is counted, 1.6σ and 1.4%; the cross-section pole mass 172.4 ± 0.7 gives 0.0934 ± 0.0009, also 1.6σ. (4) Exact agreement needs $m_t = 173.7$ GeV, 4σ above the direct measurement. (5) Re-solving Computation 176's equation for c on current inputs gives 0.9927 ± 0.0044, the same 1.6σ, with $\sqrt{\rm e}$ excluded at 149σ. Tampers: the three-loop QCD piece of the y_t matching is worth nine times the third loop of the running, so the precision lives in the matching; Computation 123's error set is identified. (6) On the 2026 Review [233] ($m_t = 172.60 \pm 0.27$, $m_h = 125.13 \pm 0.11$, α_s unchanged, $M_W = 80.3625 \pm 0.0077$) the Higgs mass's 0.07 GeV does most of the work: 0.0930 ± 0.0008 with the pole-mass ambiguity counted, 1.4σ and 1.2% (1.9σ on the quoted errors alone), exact agreement at $m_t = 173.6$ GeV, and $c = 0.9938 \pm 0.0044$. Earlier printings gave 1.6σ; the book now prints 1.4σ, on the cover too. (7) Run back down, $\lambda_{\rm SM}(M_*) = {\rm e}^{-1}/4$ gives $\lambda_{\rm SM}(m_t) = 0.1251$ against the 0.1260 the matching extracts from the observed (m_h, v), so the

prediction returns $v = 247.1$ GeV from m_h alone, 0.4% above 246.22; the earlier "0.01%" applied the tree-level relation to the running coupling at $\mu = v$ (which gives 256 GeV even on the data) and is withdrawn.

Computation 203 · §19.2.1 · [PROOF]

The single-monomial L_comm problem closed for every m, and the sharp constant certified in exact arithmetic.

Supplies the closed form Computation 157 left as the remaining work. (1) Along a fixed degree the commutator weight obeys the exact ratio test $f(a+1)/f(a) = (a+m+2)\,b/((a+1)(b+m-1))$, at least one exactly while $(m+1)b \geq (m-1)(a+1)$, and both shift weights increase with the degree, so on degree $\leq N$ the maxima sit at the top degree $s = N - 2m$: $\|T_{w^m}\|^2$ at $a = b = s/2$ and $\|[J_+, T_{w^m}]\|^2$ at $a^* = \lfloor((m+1)s - (m-1))/(2m)\rfloor + 1$; the ceiling is the sawtooth. Verified in exact rational arithmetic against every lattice point for $m \leq 8$ and $N \leq 2m + 60$. (2) The finite-N single-monomial ratio is therefore an explicit factorial ratio for every m, generalising (274); its limit is Lemma 5's closed form to 10^{-6} at $N = 4000$, and the deficit $(\mathrm{ratio}_{\mathrm{cf}} - \mathrm{ratio}_m(N))N$ tends to 1 at $m = 1$ and to 0 for $m \geq 2$, the $1/N$ term cancelling, with an $O(1/N^2)$ sawtooth left. (3) The two-monomial bridge's $1/N$ constant does not cancel: 0.85, 0.17, 0.82, 0.81 at $k = 2, 3, 4, 6$ measured to $N = 480$; extremal chains identified; a harmonic (semiclassical) correction to the chain symbol does not reproduce it. (4) Each chain norm is the top eigenvalue of a nonnegative tridiagonal matrix; Collatz–Wielandt bounds with a rational test vector enclose it exactly. Over $k \leq 20$ at the ten smallest admissible N every $G(k, N)$ is certified below 3.6736; at $(6, 12)$ the enclosure is narrower than 10^{-7} and reads $[3.67356, 3.67356]$; the widest enclosure, at $k = 18$, still certifies $G < 1.4$. Caveat: $\alpha(k)$ is an algebraic root taken at fifteen digits. Tampers: one step off a^* or one degree down lowers the weight. What remains: an analytic bound uniform in N for the two-monomial bridge. *Settled by Computation 213 and Computation 214*: the bridge's $1/N$ term is a

limit in closed form, and the sharp value is derived at its corner.

Computation 204 · §19.2.1 · [PROOF]

The $1/N$ cancellation of Computation 203 proved, and the N^{-2} sawtooth in closed form.

Both shifts move the degree by $2m$, so the degree-dependent Bergman factors cancel and $\text{ratio}_m(N)^2 = m^2 \prod_{i\le m+1}(a^*+i) \prod_{i\le m-1}(b^*+i) / \prod_{i\le m}(x+i)^2$ with $x = s/2$, a finite product. Writing $a^* = (m+1)x/m + \delta$, $b^* = (m-1)x/m - \delta$, with $\delta \in ((1-m)/(2m), (m+1)/(2m)]$ a function of s mod $2m$, the logarithm's expansion in $1/x$ is a finite Taylor expansion: the δ terms cancel at first order ($m\delta/x - m\delta/x$), the first-order sums are $m(m+2)/(2x) + m^2/(2x) = m(m+1)/x$ above and $m(m+1)/x$ below, and the second-order term is $-m(6m^2\delta^2 - 6m\delta + 1)/(6(m^2-1)x^2)$, equation (277). Verified symbolically (exact series, δ symbolic) for $m = 2$ to 12, the general-m first-order identity checked as a closed sum, and the N^{-2} coefficient verified in exact rational arithmetic at every residue of s mod $2m$ for $m \le 7$, the residual falling as $1/N$. The coefficient is $6(m\delta - \frac{1}{2})^2 - \frac{1}{2}$ times $m/(6(m^2-1))$: negative when the lattice point sits at the continuum maximiser, $3m^2/2 - 1/2$ at the ends of δ's range, so the deficit changes sign along the sawtooth. Why it cancels: first-order stationarity at an interior maximum; an $O(1)$ displacement of a^* leaves the $1/N$ term zero, a displacement proportional to x moves the limit itself below ratio_{cf}, and at $m = 1$ the maximiser is the corner $b^* = 0$, where stationarity fails and $\text{ratio}^2 = 4(N-1)/N$ keeps its $1/N$.

Computation 205 · §15.10 · [EMPIRICAL]

The M_* confrontation: three readings of the modal cutoff against the LHC, two excluded by enormous margins, the third consistent

and predicting.

The computation the book said had not been performed. *Scope, as first run*: the first measurement in the bin R3 names, ATLAS's boosted $H \to \tau\tau$ on Run 2 plus partial Run 3 [234], gives above 650 GeV a Run 2 deficit and a combined upper limit of 4.5 times the Standard Model at 95% confidence, a limit and not the 15% measurement that would exclude R3; the exposure stands untested. The readings: R1, a momentum cutoff above which no field mode exists; R2, the projection's Gaussian kernel on every vertex; R3, the kernel on the Higgs doublet alone, every vertex with a doublet leg carrying it in the momenta of its legs, gauge and fermion vertices carrying none. The data, all ATLAS at 13 TeV: inclusive jets to $p_T = 3.5$ TeV and dijets to 9 TeV with NNLO QCD [217]; dileptons above 2 TeV, 25 observed against 22 expected, highest at 4.06 TeV, contact-interaction limits of 32 to 36 TeV (constructive) and 26 to 29 TeV (destructive) [218]; $\kappa_V = 1.035 \pm 0.031$, $\kappa_F = 0.95 \pm 0.05$ [219]; $\mu_{HH} < 2.4$ [220]; boosted $H \to b\bar{b}$ limits of 115 fb above 450 GeV and 9.6 fb above 1 TeV against 18.4 and 0.13 fb [221]; same-sign WW at $2.92 \pm 0.22 \pm 0.19$ fb [222]. (1) R1 predicts zero events above 1.285 TeV: excluded at probability zero to every decimal. (2) R2 multiplies the Drell–Yan yield at 2.2 TeV by 0.053, so 12.4 expected dielectron events become 0.66 against 19 observed ($p = 10^{-20}$); linearised, it is a contact interaction with $\Lambda_{\text{eff}} = 13$ to 16 TeV by chirality against limits of 32 to 36 TeV, short by a factor 4 to 7 in $1/\Lambda^2$; the data require a universal kernel at $2.7\,M_*$ or above. (3) R3: coupling modifiers shift by -0.5 percent (1.3σ and 0.9σ from the measurements); Higgs production suppressed by $\mathrm{e}^{-p_T^2/M_*^2}$, 0.88, 0.77, 0.55 above 450, 650, 1000 GeV, invisible at present precision; $\mu_{HH} = 0.73$; with the whole doublet carrying the kernel the longitudinal cancellation is preserved, whereas a kernel on the physical h alone breaks $|a_0| < 1/2$ at 1.5 TeV; every loop of the Higgs self-energy has a doublet vertex, so the naturalness result stands, and what is given up is finiteness of loops that never touch the condensate. (4) Predictions: a 23 percent deficit above 650 GeV and 45 percent above 1 TeV in the Higgs p_T spectrum, 27 percent in Higgs pairs, modifiers within half a percent of one.

(5) Computation 10 revisited: its $M_* = 1.573$ TeV predates the correction and its $O(1)$ Wilsonian reading is R2; its structural argument is R3. Tampers: a universal kernel at $3\,M_*$ passes the dilepton limits, so the exclusion is about the scale; under R3 a kernel at $M_*/2$ would put the 650 GeV bin at 0.36 of the Standard Model, which the measurement disfavours, so R3 is falsifiable.

Computation 206 · §19.2.1 · [NUMERICAL]

The two-monomial deficit to $N = 15360$: no plateau, so the per-k constant of Computations 157 and 203 is a value, not a limit.

Before an analytic route to the bridge's $1/N$ constant is attempted, this file asks whether the number is a limit. Computation 157's chain reduction with the weights vectorised, every chain scanned at $N \leq 3840$ and a refined window around the extrapolated extremal chain above it. (1) The deficit times N rises monotonically through $N = 15360$: 0.851, 0.880, 0.913, 0.948, 0.983, 1.015 at $k = 2$ and 0.811, 0.838, 0.881, 0.934, 0.992, 1.050 at $k = 6$ for $N = 480$ to 15360, increments per doubling of 0.03 and 0.06 that shrink only slowly; run offline to $N = 61440$ the values are 1.072 and 1.162. The data do not decide between a limit approached logarithmically and a term of order $\log N/N$. (2) The commutator's extremal chain moves, its start at $0.392\,N$ to $0.420\,N$ ($k = 2$) and $0.208\,N$ to $0.222\,N$ ($k = 6$), towards but not at the single-shift maximiser's position; the tamper evaluating the norm on that maximiser's chain falls short by a growing margin. (3) The symbol's deficit moves the other way, 1.01 to 0.46 at $k = 2$ from $N = 480$ to 15360, crossing the exact value, so the symbol route to a closed form is closed. (4) Every value computed lies below the certified 3.6736 at $(6, 12)$; the uniform bound holds on every point, and whether it holds as $N \to \infty$ for fixed k is exactly what is left open, so the closure rate of Lemma 7 is established up to a possible logarithm. *Settled by Computation 213*: the value is a limit approached from below as $N^{-1/3}$, and the limit is in closed form.

COMPUTATION **207** · §20.6 · [SURROGATE]

The spectral dimension of the projected geometry: it rises between the grain and the coherence length, sits at four above d_0, and never reaches two.

The configuration-space walk of the Boolean Dirichlet form, one flip per unit of Mosco time, pushed to M by the A4 displacement covariance, has an exact return probability at resolution ϵ: a Binomial mixture over the number of odd-flipped bits of chi-square tails. Carlip's $d_s = -2\,\mathrm{d}\log P/\mathrm{d}\log s$ is read off it for $|D|$ up to 65536, with the cubic lattice $\mathbb{Z}^4$ at the same rate as the fixed-valency control. (1) The Mosco plateau is there: $|d_s - 4| < 0.1$ on a window of s whose top end grows with $|D|$ (25 to 1690 for $|D| = 256$ to 65536, the ratio per quadrupling tending to four). (2) At $\epsilon = d_0$, the resolution at which ρ distinguishes points, the curve rises from zero to a maximum of 4.46 and settles at four, an overshoot below the lattice's own 4.87: no rise. (3) At the grain $\epsilon = \ell_* = d_0/14.5$ it rises through four at $s = 2.5$ to 22.9 at $s = 13$, a few coherence lengths, and settles at four by $s = 30$; the height is a logarithm of the resolution, 8.4 per e-fold, unbounded for exact return, where the bare hypercube's $d_s = 2s$ is recovered. The mechanism is the causal set one: a first step in any of $|D|$ directions makes the return probability fall exponentially until the diffusion has spread past the resolution. (4) The lattice's overshoot is bounded and resolution-independent, so the rise is the unbounded-valency signature; reduction to two appears in neither model. Tamper: a deterministic number of flipped bits removes the not-moved term and the rise with it. Stands in for: the emergence-determined realisation, by Gaussian i.i.d. displacements with the A4 covariance; and $\mathbb{R} \times S^3$, by a flat Euclidean bulk.

Computation 208 · §18.1 · [SURROGATE]

The realised cloud as a sprinkling: PST's discreteness passes the Lorentz-invariance theorem exactly when the realised positions are independent, which A3 must supply beyond the density.

Bombelli, Henson and Sorkin [225] prove that a Poisson sprinkling selects no frame and carries no finite-valency Lorentz-invariant graph. The theorem's hypothesis is a Poisson process with the invariant volume as intensity; PST's cloud has the invariant volume as density by A2 and A3, and $|\det g|$ is unchanged by the Wick rotation, so the Riemannian and Lorentzian processes coincide. What A2's weak equidistribution does not fix is independence, and the file shows the gap on the i.i.d. surrogate in a 1+1 diamond against a lattice of the same size. (1) Both pass A2 at the same rate. (2) The index of dispersion of cell counts separates them, 0.98 against 0.00. (3) The cloud's counts in boosted blocks of equal volume keep their mean and dispersion at rapidities to 1.5; the lattice's dispersion moves from 0 to 2.2 with the boost. The cloud's causal links fill rapidity flatly over $|\eta| < 2$ on a range that widens with N; the lattice's pile up along its own null directions. (4) The cloud's links per element grow like $\log N$ (4.0 to 6.2 for $N = 250$ to 2000), the finite-valency clause and the mechanism of Computation 207. (5) $|\det g|$ agrees in both signatures. Tamper: a lattice jittered by 0.3 of its spacing passes A2 and still fails. So the theorem applies to PST exactly when the realised positions are independent: true by construction for the surrogate, and for the emergence-determined cloud one more thing P3's typicality must deliver (Computation 132: no T-channel forces the placement). *Scope, from Computation 211*: the independence is what the maximum-entropy clause means once read as a law on the realised cloud, and the rate leaves it intact at the grain. Stands in for: the emergence-determined realisation, by the i.i.d. uniform cloud of Computation 70.

Computation 209 · §18.1 · [NUMERICAL]

The classical-scaling placement of the signed-tension metric has no Euclidean geometry, so the independence of realised positions cannot be tested on it.

The natural test of Computation 208's condition is the one explicit placement the book has: classical scaling of eq. (30) on μ-typical directions $v(a)$ i.i.d. on S^6, $T(\{a,b\}) = |v(a) + v(b)|$, τ_{null} a low quantile. With $T = \sqrt{2 + 2\, v_a \cdot v_b}$ the Gram form $B = -\frac{1}{2} J d^2 J$ is, at leading order, $-(\sqrt{2}/4)$ times the centred Gram matrix of the directions: negative of rank seven. (1) At $N = 200$ to 1600 the seven most negative eigenvalues of B grow in proportion to N and match that multiple of the direction Gram eigenvalues, the eighth an order of magnitude smaller. (2) The largest positive eigenvalue grows like $\sqrt{N}$: the Euclidean side is noise and no cloud emerges. (3) The index of dispersion read on the positive-side coordinates is 0.9 and 1.2 for i.i.d. directions and 1.1 and 1.7 for a design evened on the sphere until its nearest-neighbour spread is 0.03 instead of 0.19: the test cannot tell them apart. (4) Control: with the sign of d^2 flipped the seven eigenvalues turn positive, the positions become a function of each site's own direction, and the evened design reads 0.58 against 1.27. Tamper: a cut above the bulk of the tensions reverses the signal, so the sign is set by which side of the cut the bulk sits on; eq. (30) puts it on the spacelike side, the negative case. The placement is the quantitative form of the remark in §4.12 that the Gram construction cannot supply the signature: it supplies seven negative directions, not one. The condition of Computation 208 stays P3's, with this placement ruled out as its instrument.

Computation 210 · §20.6 · [NUMERICAL]

The spectral dimension on the book's own bridged model: three between the truncation and the compact scale, zero below the truncation, no rise.

Computation 207 was computed on the surrogate. This file computes the same quantity on Computation 9's bridge: Walsh weight n mapped to degree n on S^3, eigenvalue $n(n+2)$, the binomial multiplicity absorbed into the measure so that degree n carries $(n+1)^2$, truncated at $k_n = (3 \cdot 2^D)^{1/3}$ where the mode counts match. (1) $|d_s - 3| < 0.15$ between about $5/k_n^2$ and the compact scale 0.075 for $D = 12$ to 24 ($k_n = 23$ to 369), the lower end falling as $2^{-2D/3}$; the maximum is 2.96 to 3.00. (2) Below the truncation d_s falls to zero: reduction, no rise, no two. (3) With the binomial multiplicity kept instead of reweighted, d_s rises above three to 4.5, 5.7, 6.9, 8.1 for $D = 12$ to 24: the rise of Computation 207 is carried by the Walsh modes the bridge absorbs into the measure. Tamper: without the truncation the curve is the exact S^3 one. So on the book's model the projected geometry's spectral dimension is the Mosco value above d_0 and reduction below it; the rise between the grain and d_0 is seen by a probe of the substrate's walk, not of the emergent geometry.

Computation 211 · §18.1 · [NUMERICAL]

Maximum entropy supplies the independence: the realised cloud is a sprinkling at the grain, and A3's rate touches only the scales above the constraint.

Computation 208 named independence of the realised positions as one more thing P3's typicality must deliver; Computation 209 showed the explicit placement cannot test it. This file tests the clause itself, read as a law on the realised cloud. (1) *The theorem*: among laws of N labelled positions with the same one-point marginals the product law has the largest entropy,

$H(X_1, \ldots, X_N) \leq \sum_i H(X_i)$ with equality iff the positions are independent, so maximum entropy given the density *is* independence; verified exactly on $K = 8$ cells and $N = 16$ points over $p(n) \propto \text{multinomial}(n)\, \mathrm{e}^{-J \sum_k n_k^2}$: the labelled-cloud entropy is $N \log K = 33.27$ at $J = 0$ and strictly below it on both sides, and the dispersion there is the binomial $1 - 1/K$. (2) *The constrained cloud*: the maximum-entropy law under A3's rate is the product law conditioned on the constraint set; sampled in the 1+1 diamond ($N = 4000$, the eight lowest Fourier modes held to $|\overline{f}| \leq \varepsilon$ with ε half and a tenth of the i.i.d. scale $1/\sqrt{N}$, thirty clouds each, greedy descent then Metropolis within the set), the index of dispersion is one to within 0.01 at every cell size up to 0.05 and 0.92 to 1.02 in boosted small cells at rapidities 0 to 1.5, and falls only towards the constraint's half-wavelength (cell 0.5: 0.75, 0.36, 0.24; the first is the binomial $1 - 1/4$). The constraint makes the cloud hyperuniform at the Mosco scale and leaves it Poisson at the grain. (3) The causal links of a constrained cloud carry the sprinkling's signature: 5.6 links per element, rapidities flat to within a factor 1.2 over $|\eta| < 2$. Tampers: a jittered lattice passes the same constraint six times inside the tight ε with dispersion 0.35 at the grain and boost-dependent counts (0.30, 0.38, 1.08), so the rate alone gives neither independence nor the absence of a frame; and the tight conditioning moves 150 of 4000 points and still leaves the grain at one. *What this changes*: B.7 closes on P3's footing; the fibres of ρ (B.6) are a thinning of this cloud, which preserves the Poisson property.

Computation 212 · §19.9.2 · [NUMERICAL]

What a point carries and what a device is, on the theory's own measure: fibres hold at most seven properties and are singletons almost surely, and no device of four or more properties is a substrate accident.

Two remarks shared a cause: B.6 left the fibres of ρ unsaid, and the device accident rates of Computation 156 were computed against a uniform null

because PST fixed no measure over configurations. Since Computation 141 and Computation 211 it does: directions i.i.d. under the G_2-invariant law. (1) A fibre is a clique of pairwise-null properties ($g = 0$ means $T = \tau_{\text{null}}$ exactly), and any G_2-invariant pair tension is a function of the angle alone, so when it is injective a fibre is a set of unit vectors in V_7 with one pairwise inner product c_0; its Gram matrix $(1-c_0)I+c_0J$ has rank k for $c_0 > -1/(k-1)$ and does not embed in $\mathbb{R}^7$ beyond $k = 7$. So a fibre holds at most seven properties when $c_0 > -1/6$ and at most $1+1/|c_0|$ when the null angle is more obtuse (two at the 5% quantile of the pair tensions, three at 20%, seven at the median); constructed at $k_{\max}$ for nine values of c_0 (residual 10^{-15}) and failing at $k_{\max}+1$; in $\mathbb{R}^8$ eight, so the seven is $\dim V_7$. At the Rankin bound the clique's net tension vanishes exactly: a maximal fibre is a balanced sub-configuration. (2) Under a continuous law on directions an exact coincidence $T = \tau_{\text{null}}$ has probability zero: no pair lands on it in 220 replicas up to $N = 800$, and the nearest approach falls like N^{-2}, so fibres are singletons almost surely and A1 reads $|D_n|$ as the point count. (3) Computation 156's device criterion on this measure, at two placements of τ_{null}, 100 replicas of 60 and 40 of 200: every nearest-neighbour pair qualifies, a triple in at most a few percent of replicas, nothing of four or more at either placement, where the uniform null gave 0.97, 0.29, 0.04 at sizes three to five. Tampers: a planted compact device of four is found in 48 of 50 replicas (the two misses have a background point inside the planted separation), so the zeros are the measure's and not the detector's. *What this changes*: B.6 closes, fibres being singletons μ-almost surely and at most seven in every case; the caveat on the accident rates closes the other way, the theory's own measure being sharper than the null; what stays is Computation 156's: why T has clusters at all is $T(C)$-contingent, and the bound state's dynamics is inherited.

Computation 213 · §19.2.1 · [NUMERICAL]

The two-monomial bridge constant is a limit after all: in closed form, approached from below as $N^{-1/3}$ through the Airy edge of the bridge's own chain.

Computation 206 followed the deficit times N to 15360 and saw no plateau. This file finds the constant and why it was invisible. The deficit times N is the sum of three terms, $W + Z_C - Z_T$ (an identity holding to 3×10^{-4} at every point computed): W is the weights' $1/N$ term, the two chains' symbol maxima against $2\sqrt{k}$, and Z_C, Z_T are the zero-point energies of the commutator's and the bridge's extremal chains below their symbols. (1) W converges to a closed form from the $1/s$ expansion of the Bergman weight, $w^2 = r^p(1-r)^q(1+e(p,q,r)/s)$ along a chain of degree s and slope $r = a/(a+b)$, plus a drift term for the commutator because the hopping $o_t = v_t u_{t+1}$ couples neighbouring lattice points: $W_\infty = -\sqrt{k}\,(c_C + \gamma_C - c_T) = 0.2707$ at $k = 2$, measured 0.2708 from $N = 7680$ on; without the drift term 0.129. (2) The commutator's symbol has an interior maximum in r at $r^* = 0.7143$ ($k = 2$), so its extremal chain is the one through the corner (r^*, N), starting at $(2r^* - 1)N = 0.4286\,N$, the position Computation 206's chain was drifting towards; along it the potential is quadratic, the eigenvector a harmonic ground state of width $\sim \sqrt{N}$, and $Z_C \to \sqrt{k}\,(2r^*-1)\sqrt{oS''/2}/S_C = 0.9556$, confirmed to $O(1/N)$ by moving the truncation to $2N$ so the chain has no wall. With the wall the eigenvector keeps a distance $\sim \sqrt{N \log N}$ from it and the excess decays like $N^{-1/2}\sqrt{\log N}$ (ratios per doubling 0.734, 0.732 against the law's 0.736, 0.734). (3) The bridge's diagonal chain has $r = 1/2$ fixed and a potential rising monotonically to the truncation, so its eigenvector is an Airy function pressed against the hard edge: $Z_T = A(k)\,N^{-1/3}$ with A in closed form from the first zero of Ai, 10.97 at $k = 2$ against the measured 11.26 to 11.15, falling towards it. This is the Airy edge, and it sits on the bridge's own norm rather than on the commutator's. So the deficit times N is $D_\infty(k) - A(k)N^{-1/3} + E(N)$ with $D_\infty = W_\infty + Z_{C,\infty}$: 1.2263 at $k = 2$, 1.4738 at $k = 4$, 1.5235 at $k = 6$,

approached from below, which is why no plateau was visible: at $N = 61440$ the $k = 2$ value 1.072 is still 0.15 short, and Computation 206's increments are the Airy term's decay minus the wall term's, $0.093 - 0.064 = 0.029$ at $N = 15360$. Over k: every k to 1000 closes with Computation 59's $\alpha(k)$; $D_\infty \le 2.465$ for $k \le 80$ and below the certified 3.6736 for every $k \le 482$; its envelope grows like $k^{0.22}$, crossing 3.6736 at $k = 483$, where the asymptotic regime needs N beyond $(A/D_\infty)^3 \sim 10^8$. Tampers: the drift term dropped, the Neumann Airy zero in place of the Dirichlet one (a coefficient 2.3 times too small) and a half-well zero point (three times too large) are each rejected.

Computation 214 · §19.2.1 · [NUMERICAL]

The sharp value 3.6736 **in closed form: an algebraic number at the operative corner** $(6, 12)$**, and why that corner is the maximum.**

The one quantitative thread Computation 213 left: the sharp value sits in the small-headroom regime N close to $2m$ that the asymptotic laws do not cover. Write $N = 2m + h$; the lower shift acts on $a + b \le h$ and the upper on $a + b \le h - 2$, so every chain has at most $\lfloor h/2 \rfloor + 1$ points and $G(k, 2m + h)$ is an algebraic function of (m, α, h). At the operative corner $N = 2k$ the headroom $h = 2k - 2m(k)$ grows with k; the only $k \le 80$ with $h \le 4$ are $k = 2, 3, 4$ ($h = 2$) and $5, 6$ ($h = 4$). (1) At $(6, 12)$ both extremal chains are the diagonal $(0, 0), (1, 1), (2, 2)$ (the commutator's drift towards $0.43\,N$ is a large-N effect), the weights are exact rationals, and the two 3×3 matrices have the characteristic polynomials of (275); $\alpha(6)$ is the root in $(1, 2)$ of a degree-24 integer polynomial obtained by eliminating q and $\sqrt{1 - q^2}$ from Lemma 5(c)'s closure $R_4(q) = 2\sqrt{6}$, so $G(6, 12) = 12\,(2\sqrt{6} - \sqrt{L_C/L_T}) = 3.673556413812905\ldots$ is an explicit algebraic number, inside Computation 203's enclosure and equal to the float evaluator to 10^{-13}. (2) At fixed (m, h) the map $\alpha \mapsto G$ is single-humped with the hump where the two monomials are comparable and its top rising with m (2.00, 3.30, 3.81, 4.13, 4.45 at $h = 4$ for $m = 2$ to 6); the integer lattice $(k, m(k), \alpha(k))$ samples

those curves at few points, and $k = 6$ is the only one with $h \leq 4$ and α on a hump (3.6 percent below the $(4, 4)$ top), $k = 5$ sharing that headroom with $\alpha = 28.7$ far down the tail ($G = 1.08$). Larger m gives higher humps but larger m means larger k and larger headroom. (3) Rescanned over every chain for $k \leq 80$ and $N \leq 600$ (offline; the script rescans $k \leq 40$, $N \leq 200$): the maximum of $|G|$ is 3.6736 at $(6, 12)$, then 3.084 at $(4, 10)$, 2.862 at $(9, 18)$; with Computation 213's limits (below 2.47 for $k \leq 80$) the supremum over the operative regime at $k \leq 80$ is derived at its maximiser and scanned elsewhere. Tampers: the commutator on the chain through $(2, 0)$ gives 5.58; the other real closure root 29.85 gives -5.8; α moved by 10^{-3} moves G by 3×10^{-3}. No analytic mechanism sets the value: it is the one point where a headroom of four meets a balanced bridge.

Computation 215 · §19.9.2 · [NUMERICAL]

The native sector cannot build an apparatus: Derrick's scaling on P3's functional, and the regulated lattice at PST's own discreteness, three orders below pinning.

The apparatus question's sharper half: could the theory's own field, the order parameter of P3 with the functional (1), hold a localised object at all? (1) Derrick's scaling: for a static localised configuration in three dimensions with $E_{\text{pot}} \geq 0$, $\psi(\lambda x)$ has energy $E_{\text{grad}}/\lambda + E_{\text{pot}}/\lambda^3$, strictly decreasing in λ, so every localised configuration lowers its energy by shrinking and the only static finite-energy solution is the vacuum; evaluated on a dip, a bubble of the other vacuum and a ring. (2) On the regulated lattice (Computation 200) the shrinking stops at the spacing, so a localised object could survive only by pinning, which needs the discreteness $\kappa = (m\ell)^2$ large: PST's is $(m_h/M_*)^2 = 0.0095$ with the grain and exactly 2 with the coherence length $d_0 = \sqrt{2}\,\hbar c/m_h$. Gradient flow of the lattice energy at $\kappa = 0.0095$ from four localised initial data (bubbles of radius 6 and 12, a Gaussian dip, a single flipped site) relaxes every one to the vacuum; by bisection a flipped site persists only for $\kappa > 47.0$

and a four-grain bubble only for $\kappa > 14.7$, which PST undercuts by three orders with the grain and by factors 24 and 7 with d_0. (3) The loophole named: the theory has oscillons, localised oscillations keeping half their core energy for 10^3 to 10^4 periods (a radial simulation to $t = 1000/m$); in PST's units a period is $1/m_h = 5 \times 10^{-27}$ s and the core a few $1/m_h$, a breather of the order parameter lasting 10^{-23} s, not an apparatus. Tampers: at twice the threshold the flipped site persists; with the potential's sign flipped Derrick's $E(\lambda)$ acquires a stationary point; a one-dimensional kink survives the same flow, so the theorem is dimensional. *What this changes*: "inherited" is forced, not chosen: an apparatus is a bound state of the inherited matter sector, realised in T as a cluster that (30) reads as a device (Computation 156); which clusters $T(C)$ has stays the configuration's content.

Computation 216 · §19.9.3 · [NUMERICAL]

The black-hole interior on the theory's current lengths, and the central singularity dissolved by the fibre bound.

(A) Computation 154 evaluated the validity condition $(k\,d_0)^2 \ll 1$ on Schwarzschild with $d_0 = 2.3$ nm, the Casimir bound of its day, and the chapter printed those values after d_0 moved to 2.2 am (Computation 201). Re-evaluated: $(k\,d_0)^2 = 1.9 \times 10^{-42}$ at a solar-mass horizon, $r_{\text{break}} = 37$ pm $= 1.2 \times 10^{-14}\, r_s$, enclosing 1.9×10^{-42} of the interior volume, a mass floor of 2.8×10^9 kg; on the grain ℓ_*: 9.4×10^{-45}, 6.2 pm, 1.9×10^8 kg, the floor the chapter already had. Every conclusion of Computation 154 survives by wider margins; $d_0 = 2.3$ nm reproduces every printed number (tamper), so the discrepancy was the length alone. The uncovered region holds about 10^{25} grains. (B) The initial-singularity argument the book declined to transfer is made in a different form. A point is a fibre of ρ, a clique of pairwise-null properties holding at most seven (Computation 212), so the realised content of a collapsing body, $n \gg 7$ properties, cannot occupy one point: some pair fails to be null, and its proper distance $|T - \tau_{\text{null}}|^{1/4}$ is the region's diameter.

How far from null that pair must be is a packing bound: Rankin (some pair has $\cos \geq -1/(n-1)$) and Welch (among $n > 7$ unit vectors in $\mathbb{R}^7$ some pair has $|\cos| \geq \sqrt{(n-7)/(7(n-1))}$), evaluated with Computation 212's exemplar tension at its two placements of τ_{null} beside the best configurations a descent on the sphere finds. At the median placement seven properties can be exactly null (an orthonormal frame) and eight cannot (Welch's 1/7, attained by the simplex), the deviation growing with n (0.25, 0.28, 0.39, 0.45, 0.58 at $n = 10$, 14, 20, 30, 50); at the obtuse placement only two can, Rankin attained at $n = 4$ and 7. In the proper distance a region holding n properties has diameter at least 0.56 ($n = 8$) to 0.80 ($n = 50$) at the median and 0.76 to 1.05 at the obtuse placement, in the exemplar's units, and the single-pair distance saturates at $(2 - \tau)^{1/4}$, the local character of (30). *What this changes*: the chapter's numbers move to the current lengths; the singularity as a point carrying the collapsed content is excluded on the substrate side with the strength of the fibre result and no more; what replaces it, a region whose minimal diameter the bounds give in local units and whose dynamics the projected equations do not reach, the theory does not describe. The statistical entropy, the information question and the horizon's status are untouched.

Computation 217 · §16.5 · [NUMERICAL]

What the Bernoulli measure gives the primordial perturbations, and what it does not: Gaussian as $(d_0/\lambda)^3$, white at the grain, no tilt.

The cosmology chapter carried three statements about primordial perturbations that no computation supported: Gaussianity corrections of $O(d_0^2/\lambda^2)$, a near-scale-invariant spectrum from the Bernoulli measure, and the spectrum following from the near-threshold fluctuation structure. Measured on 256^3 independent grains. (1) *Gaussianity*: a window of side λ holds $(\lambda/d_0)^3$ grains, so the central limit theorem fixes the window mean's skewness at $\gamma_1(d_0/\lambda)^{3/2}$ and its excess kurtosis at $\gamma_2(d_0/\lambda)^3$, γ_1, γ_2 the grain's own; on a skewed grain

law the fitted exponents are 1.46 and 3.18 against 3/2 and 3. The exponent is the number of spatial dimensions, not two. (2) *Spectrum*: the grain field is white, fitted slope 0.007 against zero over a decade and a half of wavenumber, and a local function of the grains (the distinction count of (14)) is white below the grain scale. White is scale-free, not scale-invariant: where Harrison–Zeldovich is $n = 1$ for the density contrast, white is $n = 0$, and Planck's $n_s = 0.9649 \pm 0.0042$ [71] is one power of k from white and 0.035 from exact scale invariance. (3) *The gap*: neither the Harrison–Zeldovich power nor the tilt is given by independence at the grain; both belong to the transfer through the near-threshold epoch, which the book does not derive. (4) *Tamper*: the same grains smoothed over four grain lengths and re-thresholded give a red spectrum (slope -4) and window statistics that Gaussianise far more slowly, so the test tells independence from correlation. The chapter's two claims are withdrawn for three statements: Gaussian at leading order with $(d_0/\lambda)^3$ corrections, white at the grain, amplitude and tilt open (§19.9.4).

Appendix C: Claims Status

This appendix consolidates every substantive claim of the book into a single structured inventory. The chapters record the status of their results in context; this appendix is the consolidated, authoritative cross-cutting index. Each item is tagged with one of eight status categories defined below.

C.1. Status categories

Postulate. An ontological input; PST has exactly three (P1, P2, P3).

Modelling choice. An item adopted within the framework where alternatives exist; documented and not derived from postulates alone.

Derived theorem. Proved from P1–P3 with all conditions in place at proof-detail level.

Theorem under modelling choice. Proved given a modelling choice; result depends on that choice.

Conditional theorem. Proved under explicit assumptions (e.g. A1–A7 of §18.1).

Identification. Structural correspondence between substrate objects and Standard-Model / general relativity (GR) / quantum field theory (QFT) counterparts.

T(C)-contingent. Configuration content; not derivable from postulates alone (structural-scope theorem of §14.3).

Open. Not settled either way: a derivation the book does not have, or a premise it carries without deriving. The register is §19.9.1 and §19.9; C.8 indexes it so that this appendix is complete.

C.1 Postulates

Postulate	**Reference**	**Statement**
P1, Property differentiation	§1.7, Chapter 2	The capacity for distinctions to exist; sole ontological primitive. Clause (iii), faithfulness (6), requires δ to record every difference the triple carries; it is a premise, not a theorem of (4) and (5) (Computation 184).
P2, Asymmetric vector tension	Ch. 3	Property differentiation carries a V_7-valued tension with $T(D) \neq 0$ generically.
P3, Modal sublimation	§4.1	G_2-invariant Landau–Ginzburg potential on V_7; sublimation instantiates the maximum-entropy (μ-typical) configuration past threshold, supplying the equidistribution condition A3 and, the typicality reaching the content of T, its non-degeneracy (Computations 193, 195); instantiated geometry follows.

C.2 Modelling choices

Three rows left this table. The direction algebra $V_7 = \mathrm{Im}(\mathbb{O})$, the minimal KO split and the maximal Dixon tensor were listed here as modelling choices in the first edition. They are not choices: the first and third are consequences of P1, a proper sub-object costing a selector P1 does not supply (Computation 168), and the second is not a selection at all, the admissible KO pair being unique (Computation 148). What remains below is a genuine choice; the V_7 decomposition behind the generation count sits in C.4 as a structural reading whose stated condition is withdrawn, not as a choice.

Choice	**Reference**	**Alternatives & rationale**
Gaussian projection kernel for Casimir ξ	Chapter 15, Computation 68	Six kernel families surveyed; Gaussian gives $\xi = 90/\pi^2$; substrate's binomial kernel converges to Gaussian (Computation 73).

C.3 Derived structural theorems (from P1–P3 at proof-detail level)

Here "proof-detail level" denotes the rigour of each derivation given its inputs, not unconditional proof: the spacetime chain below (Lorentzian signature, diffeomorphism invariance, the cosmological rows; $n = 4$ sits in C.5) rests on the equidistribution condition A3, the inequality $\eta_{K,R} < d_0^2/4$ of §18.1.6, supplied by P3's maximum-entropy sublimation (equicoercivity A7 being a theorem given A1 to A4, with A3 so supplied), and the internal-algebra and gauge rows rest on the maximal V_7 direction algebra and the maximal Dixon tensor, which earlier editions carried as extremal modelling choices in C.2 and which are consequences of P1 rather than choices beside it (§1.7, Computation 168).

Theorem	**Reference**	**Notes**
Bernoulli measure $\mu = \bigotimes_a \mathrm{Bern}(\frac{1}{2})$ as the unique measure on $\mathcal{P}(D)$ invariant under every single-property complementation	§2.5	Derived from P1's symmetry between the presence and absence of each distinction. Permutation invariance alone does *not* suffice: any weighting depending only on $\lvert C\rvert$ is permutation-invariant.
Modal threshold τ exists with specified value	§4.1	Bifurcation point of $F[\psi, \varepsilon]$ from P2 + P3.
Vacuum manifold $\mathcal{V} \cong S^6 \subset V_7$ for $\varepsilon > 0$	§4.1	G_2-orbit of vacuum direction.
Lorentzian signature $(+,-,-,-)$	Computation 2, §4.12	From hyperbolic well-posedness of Goldstone wave equation.
Diffeomorphism invariance	§4.11	From non-uniqueness of realisation map ρ.
Geroch time-orientability of emergent manifold	Ch. 5	P3-directed threshold + permutation invariance + Mosco convergence.

Theorem	**Reference**	**Notes**
Cosmological principle	§16.2	From permutation invariance of μ.
SM gauge group SU(3) × SU(2) × U(1)	Chapter 13, Computation 6	Unimodular unitaries of A_F.
Internal algebra $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$	§13.7, Computations 3, 19	From Cl(0,6) chain-algebra structure.
Born rule	§10.2	From Gleason's theorem on $L^2(\mathcal{P}(D), \mu)$.
Fermionic excitations	§11.2	Boolean-CAR isomorphism on distinction primitive.
Spin-statistics structural	§11.3	**Three closures, two of them conditional.** The substrate-derived Dirac equation closes from P1–P3. CAR convergence holds once the relations are posed slice-wise, and the slice is supplied, a slab in the time function along $\Phi(\hat{\tau})$ normalised by the proper 3-volume (Computations 186, 191, 192); W4 follows from it (Computation 185), its earlier argument withdrawn as vacuous and its substrate route closed (Computation 181). Those two hold given A1 to A4 with A3 supplied by P3 and are listed in C.5. The structural spin-statistics theorem itself is untouched, deriving statistics from Boolean exclusion without using microcausality.

Theorem	**Reference**	**Notes**
Collapse rule (measurement postulate)	§10.3	Three closures and one condition: von Neumann measurement Hamiltonian, decoherence inheritance and threshold-crossing irreversibility close; agreement with the Born rule (§10.3.2.5) does not. It holds where the tension variable is uniform in its quantile, and that is a substantive constraint on the law of T rather than a choice of scale: the comparison is cardinal, so Computation 140's monotone-reparametrisation invariance does not reach it. Computation 145 shows the condition is *satisfiable*, with the tie floor falling as $\lvert D\rvert^{-1.5}$; satisfiable was not discharged there and is supplied almost surely by P3's typicality clause (the non-degeneracy row of C.5, Computation 195). (98) does not deliver uniqueness and Born statistics at the same coupling, so the outcome separation is carried by the einselected pointer basis and the threshold supplies irreversibility.

Theorem	**Reference**	**Notes**
Substrate-side factor $Z_H(\beta_{\mathrm{KO}}) \to e^{-1}$ (substrate side of $Z^2 \approx e^{-1}$)	§12.5, Computation 100	Tensor-product factorisation forced by P1's Bernoulli product measure (substrate-side spectral-action cutoff identification), with the limiting value e^{-1} resting on the identifications $\beta_{\mathrm{KO}} = 2$ and $\Lambda^2 = D$ (C.6). The matching identification with the SM-side ratio $\lambda_{\mathrm{SM}}(M_*)/b$ (bridge premise B) is reduced within PST to a field-strength-renormalisation identity (substrate side closed; SM-side matching open) (§12.5, Computations 110–116).
Cosmological-constant sign + equation of state $w = -1$; $\Lambda_{\mathrm{PST}} = 8\pi G\rho_{\mathrm{inst}}/c^2$	§16.5	Steady-state non-dilution mechanism, conditional only on the pre-spatial framing of the substrate, which P1–P3 carry (Mechanism B). The *magnitude* of Λ_{obs} is not fixed (C.7).

Theorem	Reference	Notes
Weight profile $W(s) = -\lvert s\rvert$ of the pair sum, forced; $w = 1/3$ for the balanced pair sum	§7.4.4, Computation 158	From (31), (30) and (40) having to agree on the trace pair by pair, with no free function; the trace of the sum is the summed $T - \tau_{\text{null}}$ and vanishes in the large-window limit of a balanced cloud, so $w = 1/3$ exactly there and for every configuration shape tried; a range cutoff at d_0 is needed for convergence and nothing is read off it. Out of balance the trace is the worldline sector's kinetic scalar (Computation 159).
Dust is not in the pair sector; rest mass enters through the matter action	§7.4.6, Computation 159	The worldline sector is exactly kinetic, $w = \beta^2/3$; a photon contributes nothing from its own null worldline, which is why radiation is 1/3; the worldline share stays at a few per cent for homogeneous matter because the window $\lvert s\rvert \lesssim d_0^2$ is a neighbourhood of the light cone, and the bond kernel cannot change that, so the projection is strictly conformal there up to that share.
Polarisation count, Stefan–Boltzmann constant and the Planck spectrum from the pair sum	§7.4.6, Computations 160, 161	Equilibrium radiation of the pair sum, inside the balanced semidefinite class; the photon's dynamics were carried separately and closed by Computations 165, 166, 189.

Theorem	Reference	Notes
d^{-6} Casimir-correction power law	Chapter 15	Categorical theorem from parity + single scale d_0. Amplitude ξd_0^2 with d_0 from the normalisation premise (Computation 201): exact in form, 10^{-19} of the force at 20 nm.
Local gauge symmetry; minimal coupling	§13.3, §13.3.1	Minimal coupling D_μ follows from gauge covariance of the projected action (Computation 166). The gauged broken U(1) is massive (the Z); the photon is the unbroken Q, inherited from the SM's unbroken-generator argument.
Canonical commutation relations $[\hat{x}^\mu, \hat{p}_\nu] = i\hbar\,\delta^\mu{}_\nu$	§10.6	The Boolean gradient projects to the momentum operator, $\hat{p}_\mu = -i\hbar\,\Pi(\partial^{\mathrm{B}}_\mu)$.
Angular momentum as Noether charge; $L \in \hbar\mathbb{Z}$	§9.4, §9.5	Noether charge of the would-be Goldstone U(1) on $\mathcal{V} \cong S^1$; the integer is π_1, the unit $\hbar$ is supplied by canonical quantisation, not derived. The winding number n gives the discrete spectrum.
Uniqueness of $G_{\mu\nu}$ (Lovelock)	§8.6	Noether's second theorem plus Lovelock, given diffeomorphism invariance (derived above) and $n = 4$ (the KO split, C.5).

Theorem	**Reference**	**Notes**
Weak equivalence principle	§8.7	Single-source projection: all matter is the same projected tension, so $\nabla^\mu T_{\mu\nu} = 0$ yields one geodesic law for every test body.
Chirality of $\mathrm{SU}(2)_L$	§14.4	Structural consequence of the threshold-direction parity argument.
Horizon and flatness problems dissolved	§16.3	From the Cosmological Principle theorem above plus a constant instantiation rate j_τ; no inflationary epoch is invoked.

Theorem	Reference	Modelling choice

C.4 Theorems under modelling choice

Theorem	Reference	Modelling choice
$\xi = 90/\pi^2 \approx 9.12$ Casimir coefficient	Chapter 15, Computation 73	Gaussian projection kernel (substrate kernel converges to Gaussian at matched scaling, correction $-2/(3\lvert D\rvert)$).
$d_0 \lesssim 2.3$ nm coherence-length bound	Chapter 15	Conditional on Gaussian kernel; factor-of-$\sim$3 spread across kernel families.

Theorem	**Reference**	**Modelling choice**
Three-generation count $N_{\text{gen}} = 3$ (structural reading)	§14.2	V_7-directional decomposition into three quaternionic sub-algebras through $\hat{\tau}$; the natural structural reading. **Its stated condition is withdrawn**: this row rested on "the Dixon-synthesis block structure (Computation 48, 106)", and Computation 106's block-diagonality claim is withdrawn. Computation 175 replaces it with a sharper statement, that orthogonality is not invariance and that $\hat{\tau}^{\perp}$ is the *irreducible* fundamental **3** by Schur, so no colour-invariant generation splitting exists inside it at all. What survives is the *count*: $\hat{\tau}^{\perp}$ is 3-dimensional over $\mathbb{C}$, and that one fact is the source of both $N_c = 3$ and $N_{\text{gen}} = 3$. PST derives one three, not two, and inherits the octonionic programme's colour–generation question (Furey 2014) rather than closing it.

Theorem	Reference	Modelling choice
Colour–generation distinguishability, settled in the negative	§13.5, Computation 175	$\mathrm{SU}(3)_c = \mathrm{Stab}_{G_2}(\hat{\tau})$ acts on $\hat{\tau}^{\perp}$ as the irreducible **3**, so by Schur no colour-invariant generation splitting of $\hat{\tau}^{\perp}$ exists: colour and generation are one three, not two. What PST inherits is the octonionic programme's question of where generations live if not there; tree-level FCNC absence is not a structural theorem.

C.5 Conditional theorems (explicit assumptions listed)

Theorem	**Reference**	**Conditions**
Mosco convergence $\mathcal{S} \to M = \mathbb{R} \times S^3$	§18.1	Seven conditions A1–A7. Equicoercivity A7 is a theorem given A1 to A4, its rate condition A3 supplied by P3 (§18.1.6); A3 is the rate hypothesis, its fluctuation half settled (McDiarmid, Computation 125) and its bias half a diffeomorphism-invariant equidistribution rate $\eta_{K,R} < d_0^2/4$ supplied by P3's maximum-entropy sublimation (the realised configuration being μ-typical, its cloud equidistributed at the rate A3 states) and corroborated on the i.i.d. surrogate (Computations 70, 83); the Mosco limit thus closes under A1–A7, A1 being architectural input rather than a substrate consequence, so it is not a theorem of P1–P3 alone. This is a typicality selection internal to P3, not a fourth postulate (no P4: the apparent bulk-density freedom is a flat-chart coordinate artefact, Computation 134).

Theorem	Reference	Conditions
Spectral dimension of the projected geometry	§20, Computations 207, 210	Surrogate-based, conditional on the i.i.d. surrogate of Chapter 18 standing in for the emergence-determined realisation: four above d_0 on a window growing with $\|D\|$, at most 4.5 at the resolution d_0, rising to about 23 at the grain before settling (Computation 207); on the book's own bridged model, three between the truncation and the compact scale with no rise (Computation 210). With the full causal set and against CDT, as a computed statement; no reduction to two appears in either model.
Dirac-operator convergence (relaxed criterion)	§18.3, Computations 9, 65, 150, 213, 214	Relaxed-criterion closure delivered by the L_round/L_comm lemma chain; the uniform bound on $k \leq k_D$ is derived at its maximiser (the sharp value 3.6736 an explicit algebraic number, Computation 214) and in its large-N limit (closed form, Computation 213), scanned between; a proof over the whole regime remains its one quantitative refinement (existence of a finite bound there is analytic, as is its $D^{-1/4}$ mechanism; Ch. 19).

Theorem	Reference	Conditions
Slice-wise CAR limit and microcausality W4	§11.3, Computations 185, 186, 191, 192	A1 to A4 with A3's rate supplied by P3, the standing of equicoercivity A7, and the Lorentzian signature of §4.12. No site is realised on a measure-zero slice, so the relations are posed on a slab normalised by the proper 3-volume of the slice (Computation 186); the slab limit holds for every cloud whose discrepancy vanishes faster than the thickness, by Koksma–Hlawka, with no independence (Computation 191); the time function is cosmic time along $\Phi(\hat{\tau})$, its foliation forced by isotropy, its level sets the compact Cauchy surfaces, invariant under relabelling (Computation 192); and W4 is the corollary, S being the causal propagator on a globally hyperbolic spacetime (Computation 185, [213, 214]). The substrate-correlator route to W4 stays closed (Computation 181).

Theorem	Reference	Conditions
Orbital angular momentum, $L_z = nQ$	§9.5, Computation 194	The tensor from the Goldstone action and the integer from $\pi_1(S^1)$ are derived, and $L_z = nQ$ is an identity for every axisymmetric configuration, Q the Noether charge. The unit is conditional on the absolute normalisation of the Goldstone stiffness r_0^2, the canonical-normalisation premise of C.8: a derived unit is $r_0^2 = \hbar c/(\kappa d_0^2)$, with $r_0 = v$ and the derived d_0 giving $\kappa = m_h^2/(2v^2) = 0.13$ (Computation 201; 1.2×10^{-19} at the Casimir bound, Computation 197), and $\hbar$ is inherited at (106). The far field is the global vortex's, L_z growing as R^3 and the energy logarithmically, and $\pi_2(S^1) = 0$ excludes the spherical defect.

Theorem	**Reference**	**Conditions**
Non-degeneracy of T's contingent content: Born agreement on the law of T, and non-balance	§10.3.2.5, §3.5, Computations 190, 193, 195	P3's typicality clause read as reaching the contingent content of T (§1.7), the footing A3 stands on. The graph half of Born agreement is a theorem (Computation 190); the two faces are one open dense condition, the magnitudes off the tie set and the directions off the balance set (Computation 193); the maximum-entropy law on the content has a density, so both faces hold almost surely under it, $\mu_{\pm}$ being its antipode (Computation 195). The reference measure is load-bearing: quantised directions would balance.
Only the magnitude ordering of T and $\hat{\tau}$ are load-bearing	§4.1, Computations 126, 140, 195	Established for every derivation but the Born comparison by the reads the body makes (Computation 140), and for that one cardinal read by the row above; the residual per-configuration content is contingent (§14.3) and a bounded underdetermination, not an open derivation.
Stress-energy = projected tension; EFE as Seeley–DeWitt term	§7.8, Computation 5	Standard CC spectral-action machinery applied to substrate triple.

Theorem	**Reference**	**Conditions**
Newton's G as consistency relation among (m_h, M_P, v)	§8.5	Conditional on heat-kernel CC matching.
$m_h^2 = (3/2)M_*^2/(16\pi^2)$ scalar self-mass; $M_* = 4\pi m_h\sqrt{2/3} \approx 1.285$ TeV	§12.2, Computations 93, 97, 109	Parameter-free one-loop result *for the Higgs self-loop in isolation.* The full-SM extension (separately accounting for the EFT-layer top/W/Z quadratic sensitivities) does not exist within PST (Computation 120): the Veltman residual at M_* is large and negative, $\mathcal{V}(M_*) \approx -3.3$ (§12.2), with no zero crossing, so M_* is irreducibly a scalar-sector result, not a full-theory prediction. Confronted with the LHC in Computation 205: survives in the Higgs-doublet reading of the cutoff, which is §15.2's, and predicts a 23 percent deficit in Higgs production above 650 GeV of transverse momentum; the literal and the universal-kernel readings are excluded by the jet and dilepton spectra.

Theorem	**Reference**	**Conditions**
$Z^2 \approx e^{-1}$ field-normalisation identity / Bridge premise (B)	§12.5, Computations 100, 110–116	Substrate-side factor $Z_H(\beta_{\mathrm{KO}}) \to e^{-1}$ derived from P1–P3 (C.3, Computation 100). Matching identification with $\lambda_{\mathrm{SM}}(M_*)/b$ (bridge premise B) is REDUCED to a field-strength-renormalisation identity, substrate side closed in the substrate limit: a wave-function (field-strength) renormalisation, multiplicative rather than the additive Wilsonian matching the measure-Jacobian route excludes (Computation 103), forced by P1's product structure into the substrate-factor form (Computation 111), substrate-side value the embedding norm of the symmetric Higgs vacuum (Computation 114), with the total reduction supplied by a unique-soft-mode theorem (Computations 115, 116).

Theorem	Reference	Conditions
(the Z^2 row, continued)		Open, indexed in C.8 under Bridge premise (B): the matching of the substrate factor to the SM-measured coupling, resting on the emergence premise F4 (§12.6.2.1) plus the RGE test (Wilsonian threshold matching, quantified by a three-loop error budget as a rigid, D-independent prediction $\mathrm{e}^{-1}/4 = 0.09197$ consistent with the SM-run 0.0930 ± 0.0008 at 1.4σ on current inputs, falsifiable as the top mass tightens, Computation 202; Computation 123's 1.0σ rested on an anchor the literature's formulas do not reproduce). Substrate-side residual: only the standard $O(1/D)$ caveat on the value e^{-1}; the all-orders decoupling is closed exactly and non-perturbatively by sufficiency of the empirical mean (Computation 121), the Morse–Bott lemma (Computation 118) being its shadow, and F8's linearity controls only that $O(1/D)$ (Computation 122).

Theorem	**Reference**	**Conditions**
Four-dimensional spacetime $n = 4$	§1.5	From the minimal Connes-spectral-triple split given the KO-6 substrate; conditional on the spectral-triple framework assumption, which C.8 carries.
Weak-field limit: Poisson's equation recovered	§8.2	Given the projected field equations (the Einstein-equation rows above); G enters as the unit-conversion factor.
Friedmann equations; accelerating solution	§16.6	Given the projected field equations and the ongoing-sublimation source.
Conservation of $\Pi_{\mu\nu}$; the divergence-free ambiguity	§7.4.4, Computations 162, 163	Given the identification, exact on balanced clouds by (46), the traceless part of a conserved tensor being conserved when the trace is constant, verified on a Friedmann background with a massive scalar (Computation 163); without it, the residual is bounded and falls with the pair count (Computation 162). The Irving–Kirkwood ambiguity is restricted by tracelessness to a global harmonic function, which changes no integrated quantity.

Theorem	Reference	Conditions
The interacting vacuum of the regulated scalar sector	§12.2, Computations 198, 200	Conditional on the cutoff being kept, which is the theory's own statement of it: in finite volume the Hamiltonian is a Schrödinger operator with a quartic potential, self-adjoint with discrete spectrum and a unique strictly positive ground state; the infinite-volume Euclidean measure exists by Griffiths monotonicity with Dirichlet data and the quartic's uniform bound; reflection positivity returns a Hamiltonian bounded below. Stability at two loops (Computation 198) is an expansion about it. The gauge and fermion sector below M_* is the cutoff Standard Model's own existence problem (the weak-coupling infinite-volume limit, chiral fermions on a regulator), inherited rather than introduced.
Coherence length $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am	Chapter 15, Computation 201	Conditional on canonical normalisation: the ξ_ψ of (205), $\sqrt{c_{LG}/(2\lvert a_{LG}\rvert)}$ with $\lvert a_{LG}\rvert = m_h^2/4$, the frame (neutral or doublet, half-coefficient or not) fixing an order-one factor; $d_0/\ell_* = \sqrt{2}\,M_*/m_h = 14.5$; nine orders below the Casimir bound; the Casimir correction 10^{-19} of the force at 20 nm.

Theorem	**Reference**	**Conditions**
Gauge anomaly cancellation forces $N_c = 3$	§13.8, Computation 136	*Conditional on the Standard-Model hypercharge assignments.* Explicitly *not* a derivation of anomaly cancellation from PST, but a consistency check: all six triangle coefficients vanish iff $N_c = 3$, the dimension PST fixes independently as the $M_3(\mathbb{C})$ colour factor.

Identification	Reference	Content

C.6 Identifications

Identification	Reference	Content
Substrate modal field = SM Higgs at different scale	§15.2	$\phi_{\text{Higgs}} = \Pi(\psi)$, $m_h^2 = 4\varepsilon$.
Substrate Higgs Hamiltonian $H_{\text{Higgs}} = \bar{X}$	Computation 89, §12.5	From additivity + matched scaling + uniformity; no free parameters.
Matched-scaling exponent $\beta_{\text{KO}} = 2$	Computation 100, 102	Spectral range of a single binary distinction, $2 = 1 - (-1)$, forced by P1's binary alphabet (equivalently the one-bit Clifford $\text{Cl}(1,0)$ range). The KO Bott-periodicity value $\text{KO}_{\text{total}} \bmod 8 = 2$ shares the value but is not the load-bearing source.
Cutoff identification $\Lambda^2 = D$ (the spectrum's matched scale)	Computation 100	Substrate dimension as natural matched scale for Boolean Laplacian's spectrum.
$\text{U}(1)_Y$ hypercharges	§13.9	A *determination*, not a derivation: four structural facts fix all six hypercharges with no free parameter beyond the $\text{U}(1)_Y$ normalisation, but the Gell-Mann–Nishijima relation and the empirical nucleon charges enter as *inputs*, not PST outputs.

Identification	**Reference**	**Content**
Conformal part of $\Pi_{\mu\nu}$ = conformal part of $T_{\mu\nu}^{(\text{fields})}$	§7.5, Computations 158, 164, 196, 199	The conformal part of the pair sum reaches exactly the metric-traceless tensors obeying the null energy condition (Computation 199), so every classical field's conformal part is reachable, static electric and magnetic fields included; on balanced clouds the conformal part is the pair sum itself (Computation 158); the magnitude is the state's radiation content, the truncation spacing being comoving (Computation 196). An identification, not a derivation of $T_{\mu\nu}^{(\text{fields})}$ from the pair sum; the trace of an unbalanced sum, the worldline sector's kinetic scalar (Computation 159), is not used.

C.7 T(C)-contingent (configuration content, not derivable from postulates alone)

Item	Reference	Status
Yukawa mass hierarchy (quarks + charged leptons)	§14.3, Computations 76–80	T(C)-contingent; structural-scope theorem (Computation 1) forces S_3 symmetry at leading order.
CKM mixing matrix	§14.3	T(C)-contingent; same structural-scope theorem.
Neutrino masses + PMNS matrix	§14.3	T(C)-contingent.
Λ_{obs} cosmological-constant magnitude	§19.9, Computations 82, 188	T(C)-contingent. Sign and equation of state $w = -1$ remain structural theorems; the theory's own area-law and finiteness resources do not reach the magnitude (Computation 188).
Clustering of the tension structure, i.e. that a laboratory device exists at all	§19.9.2, Computation 156	T(C)-contingent. The *compactness* of a device is a corollary of (30), and being one is an open condition with sharp margin $m/2$; why T clusters is configuration content.

C.8 Open (not settled either way)

This section indexes §19.9.1 and §19.9; those sections state each item in full and this table exists so that no substantive claim of the book is absent from this appendix.

Item	Reference	Status
Canonical normalisation of the emergent field (P3's gradient term)	§12.2, Computations 173, 174, 176, 177, 197, 201	**Carried as a named premise, not derived.** Reading A of §12.2: $\lambda_{\mathrm{PST}} = 1/4$ is the canonical quartic, the rival reading excluded at 121σ (Computation 176) and 9.5σ (Computation 173) with the regulator escape closed (Computation 174); the Mosco limit does not fix it (Computation 177) and the derived unit of angular momentum does not either (Computation 197). It fixes the coherence length, $d_0 = \sqrt{2}\,\hbar c/m_h = 2.2$ am (Computation 201), and with it the Casimir amplitude. The Higgs agreement (1.4σ on current inputs, Computation 202) is its check, not its derivation.
Bridge premise (B)	§19.7	Reduced, not closed: the SM-side matching rests on F4.
Framework premise F4	§12.5	Carried as an explicit unproved premise.
Matched-scaling condition A1	§18.1, Computations 179, 182	Architectural input, not a substrate consequence. The Mosco chain rests on it, and so, since Computation 182, does $\Pi_{\mu\nu}$'s magnitude: A1 is what names the realisation-point spacing on which that magnitude (C.6) rests.

Item	Reference	Status
Connes spectral-triple framework assumption	§1.5	Carried as a framework assumption; $n = 4$ is enforced by it.
Odd-$\lvert D\rvert$ parity selection	§1.5, Computation 3	Carried as a named premise: KO-6 holds on the odd-$\lvert D\rvert$ sub-limit, the one underived substrate choice behind the minimal split (Computation 148).
Residue: localising an apparatus in the property domain	§19.9.2, Computation 215	Not a derivation owed: an apparatus is inherited matter by necessity, and which clusters $T(C)$ has is configuration content (the clustering row of C.7).
Residue: the black-hole interior	§19.9.3, Computations 154, 216	Horizon, light and Wald entropy closed; the projected equations stop at $r_{\text{break}} = 37$ pm, where a point carries at most seven properties and the singularity has no substrate counterpart; what that region does, the theory does not say.
Residue: the uniform bound beyond $k = 482$	§19.9.4, Computations 213, 214	The sharp value 3.6736 is derived at its maximiser and the limit law holds for every $k \leq 482$ (the closed part is in the C.5 Dirac-convergence row); a proof over the whole regime and the growth of $D_\infty(k)$ beyond it are not supplied.

Appendix D: Comparison to Other Theories

This appendix tabulates the substrate-style and emergent-physics theories that share posture with PST: each posits a layer beneath spacetime, causality, or the Standard Model from which the macroscopic structure is supposed to follow. The narrative comparison sits in Chapter 17; the table below condenses each theory along five axes for side-by-side reading.

Inclusion criterion. A theory is included if it either (i) posits a pre-geometric primitive from which spacetime or causality is derived, or (ii) derives a Standard-Model-relevant structure from an underlying programme. Theories PST recovers as effective descriptions (General Relativity, Quantum Mechanics, Quantum Field Theory) are not in the comparison set; for those see Chapter 17.

Axes. The columns are: the theory with its attribution; the substrate primitive (what is taken as fundamental); the spacetime-emergence mechanism (how, if at all, spacetime is derived); Standard-Model coverage (whether gauge group, matter content, or families are addressed); and the relation to PST (whether PST agrees, extends, contrasts, or sits orthogonal). Entries are deliberately compressed; the cited references, where given, supply the full treatment.

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
Reference row.				
PST (Meelker, 2026)	Property differentiation with V_7-valued directional content (D, δ, v)	Mosco convergence of V_7-Boolean Dirichlet form to Laplace–Beltrami on $\mathbb{R} \times S^3$ along $\hat{\tau}$	$\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ derived; one generation from $\mathrm{Cl}(0,6)$ minimal ideal; $N_{\mathrm{gen}} = 3$ from quaternionic sub-algebras of $\mathbb{O}$ containing $\hat{\tau}$	
Pre-geometric / combinatorial substrate.				

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
Wolfram Physics Project (Wolfram, Gorard, 2020+) [123]	Spatial hypergraph with local rewrite rules	Causal graph from update sequence; (3+1)D from rule-space exploration	Programmatic; gauge groups not yet derived from rules	Closest in posture (no presupposed geometry; everything from substrate updates); differs on substrate nature (computation vs. distinction-plus-direction); PST derives SM from internal algebra rather than searching rule space
't Hooft CA Interpretation ('t Hooft, 2014) [124]	Deterministic Planck-scale cellular automaton	Spacetime lattice presupposed; QM emerges from CA bookkeeping	Programmatic; SM not derived	Shares "QM is derived" thesis; PST's Mosco-convergence route delivers CCR from substrate without requiring determinism
Group Field Theory (Oriti, Freidel, Pereira, 2000+) [125]	Second-quantised spin-network vertices on a Lie group	Spacetime as condensate phase of GFT (Bose–Einstein-like)	Silent on matter content	Combinatorial substrate analogue of PST's algebraic substrate; both invoke limit theorems; GFT yields LQG-style quantum geometry, PST yields a Connes spectral triple

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
Causal Set Theory (Bombelli et al., 1987+) [4, 5, 232]	Locally finite partially ordered set of events	Spacetime continuum as dense Lorentz-invariant "sprinkling" limit; order plus number fixes the metric [223]	Scalar fields carried on the order [226]; gauge structure and fermions not derived	Pre-geometric combinatorial substrate; PST adds V_7-valued directional content beyond bare order, supplying the gauge structure the order lacks. Three points of contact: everpresent Λ against PST's exact $w = -1$ (§16.9); the spectral dimension, which rises at short scales on the full causal set as PST's does between the grain and d_0 (§20.6, Computation 207); and the sprinkling theorem [225], that no finite-valency graph is Lorentz invariant. PST's discreteness passes the theorem exactly when the realised positions are independent, Poisson-type, which A2's weak equidistribution does not fix (a lattice satisfies it and fails) and which P3's maximum-entropy clause supplies along with the density: given the density, maximum entropy is independence, and the rate A3 states conditions only the modes above the Mosco scale, so the conditioned cloud keeps dispersion one under boosts at the grain (Computation 208, Computation 211); the explicit scaling placement cannot test it, having no Euclidean side (Computation 209)

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
Loop Quantum Gravity (Rovelli, Smolin, 1986+) [6]	Spin-network states quantising SU(2) holonomies and fluxes	Continuum spacetime as semiclassical limit of spin-foam histories	Not derived; matter coupled by hand	Spin-network combinatorial substrate; PST's substrate is algebraic (Walsh modes / Cl(0,6)) rather than spin-network; both aim at background-independent quantum geometry
Penrose Twistor Theory (Penrose, 1967+) [126]	Twistor space $\mathbb{T} \cong \mathbb{C}^4$ encoding null geodesics + helicity	Spacetime points as projective null lines in twistor space (Klein correspondence)	Conformal kinematics; gauge structure later via twistor-string/ amplituhedron	Pre-spacetime spinor programme; PST's distinction substrate is logically prior to both spacetime and the twistor encoding
Energetic Causal Sets (Cortês–Smolin, 2014) [127]	Events with energy–momentum + causal poset	Spacetime as derived from energetic causal poset	Programmatic	Extends causal sets with dynamical input; PST has richer substrate (V_7 directionality + Bernoulli measure) and a spectral-triple emergence route
Hardy Causaloids (Hardy, 2005+) [128]	Causaloid: operational data structure encoding probabilistic causal correlations	Causal structure primitive; geometry derived	Operational, silent on specific gauge/matter	Operational pre-geometric framework; PST is an ontological pre-geometric framework with explicit derivation chain
Causal Fermion Systems (Finster, 2010+) [129]	Measure on operators of finite spatial rank in a Hilbert space (no a priori spacetime)	Spacetime + causal structure emerge from operator-support overlap; Lorentzian signature from causal action principle	Derives Dirac sea + a Standard-Model-like gauge structure (chiral fermions, weak isospin) from the causal action	Closest mathematical kin: both use Hilbert-space-with-extra-structure as substrate; CFS uses measure theory + causal action, PST uses Connes spectral triples + Mosco convergence; CFS gauge derivation runs through a different mechanism
BFSS Matrix Models (Banks–Fischler–Shenker–Susskind, 1996) [130]	Large-N matrices of $D0$-brane positions and spinor partners	11-D supergravity as the $N \to \infty$, 't Hooft limit; spacetime as eigenvalue distribution	Low-energy SUGRA spectrum	Matrix-model substrate; PST is spectral-triple substrate; both invoke large-N/coarse-graining limits but in mathematically distinct categories

Theory	**Substrate primitive**	**Spacetime emergence**	**Standard-Model coverage**	**Relation to PST**
General Substrate Theory (GST) (informal proposal; no published source to cite)	Causal pacing in place of geometry; physical laws as coherence-conserving structural adaptations	Spacetime replaced by causal-pacing relations; relativistic and quantum-discreteness phenomena recovered from coherence-conservation constraints	Programmatic; not specific on gauge/matter content	Shares pre-geometric posture; PST commits to a specific substrate (V_7-directional distinguishability) and a definite emergence theorem (Mosco convergence) rather than an adaptive coherence-pacing principle
"Linkon" pre-geometric substrate proposal (informal proposal; no published source to cite)	"Linkons": primitive relations that self-organise	Continuous information fields and gravity emerge from coherence among linkons rather than from a pre-existing spatial manifold	Programmatic	Closest in stated spirit, "relations as primitive" parallels PST's distinguishability; PST commits to specific algebraic structure (V_7, Bernoulli measure, Landau–Ginzburg potential) and a Connes spectral-triple realisation, where that proposal as described remains schematic
Substrate Theory of Everything (Reed, 2024–2026) [143]	Single scalar substrate field with Mexican-hat potential and U(1) gauge structure; eight ontological postulates; two-frame ontology distinguishing undisturbed substrate (F_1) from perturbed observable form (F_2)	Spacetime emerges from substrate perturbations; Newton, special relativity, general relativity recovered as theorems	Standard-Model gauge structure and particle-mass cascade derived from U(1) cascade; dark matter / dark energy and matter–antimatter asymmetry addressed; testable predictions at HL-LHC, Gaia DR4, CMB B-modes, LISA	Most parallel in posture: another substrate from which spacetime, gravity, and SM physics are derived, with a Mexican-hat-style potential at the centre. Differences are technical: Reed uses a scalar substrate with a single U(1) producing the SM gauge cascade, where PST uses a V_7-directional distinguishability substrate, G_2-invariant LG potential on V_7, and the Connes spectral-triple route; PST's gauge group emerges as the unimodular unitaries of $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ (itself derived from P1–P3) rather than from a single U(1) cascade

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
(continued: 2026 release)			Adds a proposed resolution of the cosmological-constant problem, reducing the quantum-field-theory discrepancy from ∼122 to ∼1 orders of magnitude via a holographic bound on F_2 substrate displacement	On the cosmological constant the two part company on what is being claimed: Reed offers a numerical resolution of the vacuum-energy discrepancy, where PST fixes the location, sign and equation of state $w = -1$ of the accelerating component from ongoing sublimation and leaves the static valley depth to renormalise like any vacuum energy, relocating and disarming the 120-orders mismatch rather than resolving it numerically (§9.1)
Replication-Based Pre-Geometric Cosmology (Montanari, 2025) [142]	Finite "unstructured energy" whose sole intrinsic property is vibration, supported on an open bounded spectral interval $(\omega_{\min}, \omega_{\max})$ (finiteness forbids eternally stable configurations); evolved by competing excitation vs. randomisation in an energy-conserving Langevin/Fokker–Planck regime, yielding a metastable unstructured state	Not derived in Paper I (deferred to Papers II/III). Only pre-geometric precursors arise: emergent time as maximal-correlation ordering $C(t,t') = \int a(\omega,t)a(\omega,t')\,d\omega$ plus a discrete replication tick $\Delta\tau$; an acyclic pre-causal partial order (light-cone analogue); and a Nyquist minimum resolution $\sim 1/\omega_{\max}$ from spectral compactness	None; purely cosmological, with no gauge group, matter content, or families addressed (axioms PG1–PG9 + qualitative dynamics only). Persistent structure arises solely via a $1 \to n$ replication instability of coherent spectral units (eigenpatterns of a universal excitation drift), consuming finite substrate energy and terminating on depletion	Shares PST's pre-geometric posture, finiteness as the load-bearing constraint, emergence of time and causality from substrate dynamics, and a competing order/disorder instability echoing PST's tension-threshold (G_2-invariant Landau–Ginzburg) sublimation. But Montanari is cosmology-only and outline-level (axioms; no concrete spacetime derivation or predictions yet), whereas PST carries the same posture through to a derived $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$, three generations, and falsifiable numbers (a d^{-6} Casimir correction at $d_0 = 2.2\,\mathrm{am}$, bounded by data at 2.3 nm; $\lambda_{\mathrm{SM}}(M_*) = \mathrm{e}^{-1}/4$ at 1.4σ)
Emergent-gravity programme: gravity as effective description of underlying microphysics.				

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
Sakharov Induced Gravity (Sakharov, 1967) [131]	Standard-Model quantum-field vacuum fluctuations	Einstein–Hilbert action as one-loop effective action of matter fields	Presupposes SM matter	Historical precursor of all emergent-gravity proposals; PST replaces "vacuum fluctuations" with a precausal substrate and derives the SM rather than presupposing it
Thermodynamics of Spacetime (Jacobson, 1995) [19]	Local Rindler-horizon degrees of freedom (unspecified)	Einstein equation derived as the Clausius equation of state $\delta Q = T\mathrm{d}S$ on local horizons	Not addressed	Parallel "gravity from underlying microphysics" claim; PST specifies the microphysics (V_7-substrate) and recovers the same Einstein equation via the spectral action
Atoms of Spacetime (Padmanabhan, 2002+) [132]	Microscopic "atomic" degrees of freedom of spacetime horizons	Gravity as horizon thermodynamics; equipartition over the surface degrees of freedom yields Einstein equations	Not addressed	Same emergent-gravity claim as Jacobson; PST's route is spectral-action (matter + gravity coupled at leading order) rather than horizon-thermodynamic
Entropic / Emergent Gravity (Verlinde, 2010, 2016) [18, 133]	Holographic information / quantum entanglement between regions	Gravity as entropic force from holographic counting; modified inertia in low-acceleration regime	Dark-matter phenomenology suggested; SM not derived	Emergent-gravity programme; PST's gravity emerges from spectral action on the substrate, not from entanglement entropy bookkeeping
Quantum-information substrate: spacetime from entanglement.				
ER=EPR / It from Qubit (Maldacena–Susskind 2013; Van Raamsdonk 2010) [134, 135]	Entanglement structure of the boundary quantum theory	Bulk geometry from entanglement entropy (Ryu–Takayanagi); ER bridges as EPR pairs	Inherits boundary CFT's matter content; not derived	Spacetime-from-entanglement vs. PST's spacetime-from-Boolean-Dirichlet-Mosco; complementary programmes; PST gives a substrate finite in internal structure and an explicit limit theorem, IfQ relies on holography

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
Tensor-Network Emergence (Swingle 2009; Vidal MERA; see also Kitaev–Preskill on topological entanglement entropy) [136, 137, 107]	Hierarchy of tensor-network entanglement (MERA / HaPPY codes)	AdS geometry from MERA scale layers; bulk locality from network connectivity	Silent on matter content	Operational realisation of It-from-Qubit; geometry from QI structure rather than from a logical primitive
Algorithmic Information Substrate (informal proposals; no published source to cite)	Algorithmic information / computational complexity of the substrate's update history	Quantum vs. classical regime separated by a critical algorithmic-complexity threshold ($\sim$ 50 bits in some proposals): below threshold the substrate maintains quantum superposition; above it, classical irreversibility	Not derived; treated as effective consequence of complexity-threshold crossing	Complementary in mechanism: PST yields the quantum/classical transition via Mosco convergence of Boolean Dirichlet forms and the modal threshold τ rather than via an algorithmic-complexity cut-off; both share a discreteness $\to$ continuum scaling structure
Foundational / philosophical substrate proposals.				

Theory	Substrate primitive	Spacetime emergence	Standard-Model coverage	Relation to PST
Octonion-Substrate Programme (Dixon 1994; Manogue–Dray 1999+; Furey 2014+) [138, 139, 121, 122]	Division-algebra chain (e.g. $\mathbb{C} \otimes \mathbb{H} \otimes \mathbb{O}$); octonions as fundamental	Spacetime not derived; programme is internal-algebra-focused	One generation of fermions from minimal ideal of Cl(6); SU(3) × SU(2) × U(1) from unimodular unitaries	Closest *technical* kin: PST extends Furey 2014 by deriving $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$ from P1–P3 (rather than positing the chain algebra) and pairs the internal algebra with explicit spacetime emergence; PST supplies a different *selector* for the generation count, reading $N_{\text{gen}} = 3$ off the quaternionic sub-algebras of $\mathbb{O}$ that contain the modal-sublimation direction $\hat{\tau}$ rather than off the algebra alone; the programme reaches three generations by its own route (Furey 2018), and PST inherits its colour–generation distinguishability caveat rather than removing it (C.4)
Constructor Theory (Deutsch & Marletto, 2013+) [140]	Counterfactual possible/impossible transformations	Spacetime presupposed; not derived	Not addressed	Shares "logical primitive" stance with PST's distinguishability postulate; orthogonal mathematical machinery (possibility/impossibility lattices vs. Connes spectral triples)
Spencer-Brown Distinction Ontologies (Spencer-Brown, 1969; Kauffman; Varela) [61]	The act of making a distinction	Not formalised; mostly philosophical	Silent	PST is the closest *formalisation* of "distinction as primitive" with an explicit derivation chain to physics; shares the foundational starting point but commits to additional structure (V_7 directionality, Bernoulli measure, Landau–Ginzburg potential)

Theory	**Substrate primitive**	**Spacetime emergence**	**Standard-Model coverage**	**Relation to PST**
Wheeler's "It from Bit" / MUH (Wheeler 1989; Tegmark 1998+) [62, 63]	Binary informational acts (Wheeler) / all mathematical structures (Tegmark)	Programmatic; no specific derivation	Programmatic	Programmatic ancestor; PST instantiates the philosophy with concrete substrate + derivation chain rather than leaving emergence as a slogan
Information Substrate Theory (Kriger, 2026) [141]	"Information" as the name for *differentiatedness* in being; a single axis-free substrate of which Shannon entropy, Kolmogorov complexity, thermodynamic entropy etc. are specific projections (along time, space, energy, ...)	Spacetime not derived; affirmed as one projection / coordinate reading of the same substrate	Affirms existing SM; not derived	Closest in *ontology*: both PST and IST take differentiatedness as the foundational primitive (PST's "property differentiation" ↔ IST's "information as the name of differentiatedness"). PST commits to a specific algebraic realisation (V_7-directional distinctions, Bernoulli measure, G_2-invariant Landau–Ginzburg potential, Connes spectral triple) and an explicit derivation of spacetime and the Standard Model; IST keeps the substrate axis-free and reads existing physics off it as projections rather than deriving the projections from postulates

D.1. Conclusion of comparison

The table holds twenty-seven entries, PST and twenty-six others. PST is the only entry in the table that does all five of the following at once.

(i) It posits a pre-geometric substrate finite in internal structure, not a finite universe.

(ii) It derives spacetime as a limit theorem from it.

(iii) It derives the structural ingredients of the $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ gauge group from the three postulates.

(iv) It places $N_{\mathrm{gen}} = 3$ in Furey's six-SU(3)-triplet decomposition of $\mathrm{Cl}(0, 6)$ as the natural reading.

(v) It delivers a specific Standard-Model coupling value, $\lambda_{\mathrm{SM}}(M_*) = b \cdot \mathrm{e}^{-1} = \mathrm{e}^{-1}/4 = 0.0920$, where $b = 1/4$ is the Landau–Ginzburg modal quartic supplied by P3, matching the observed value at 1.2%. The substrate-side factor e^{-1} is derived from the postulates; the SM-side matching is reduced within PST to a field-strength-renormalisation identity resting on the emergence premise F4 (§12.6.2.1) plus the renormalisation-group equation [RGE] test, so it is not parameter-free (§19.7).

D.2. The two nearest entries

Two rows deserve a word, because a quick reading of them might suggest they come closer to the five axes than they do.

Substrate Theory of Everything (Reed, 2024–2026) [143] reaches further into the Standard Model than any other entry in the table, and its posture is the most parallel to PST's. The route, however, is a different one. Reed's substrate is a continuous scalar field that already carries a U(1) gauge structure and a Mexican-hat potential, laid down in eight ontological postulates; the Standard Model is then reached by a cascade from that U(1). PST's substrate is finite in internal structure and carries no gauge group at all: the gauge group appears only at the end of the chain, as the unimodular unitaries of $A_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$, and that algebra is itself derived from P1–P3 rather

than assumed. On axis (i), then, Reed's substrate is not finite in internal structure; on axis (iii), its gauge content is posited rather than derived.

The divergence is sharper than a difference of machinery, and it runs in both directions. Reed derives a particle-mass cascade, which PST does not claim and, by the structural-scope theorem of §14.3, cannot claim: the fermion mass and mixing spectrum sits in the contingent configuration content $T(C)$, not in the layer fixed by the postulates. The 2026 release goes further still and offers a numerical resolution of the cosmological-constant problem, reducing the quantum-field-theory discrepancy from some 122 orders of magnitude to about one via a holographic bound on substrate displacement. PST does not claim that either. What PST fixes structurally is the location of the vacuum, the sign, and the equation of state $w = -1$ of the accelerating component, all carried by ongoing sublimation. The magnitude of the static valley depth renormalises like any vacuum energy and is not derived. On PST's own account, then, the 120-orders mismatch is relocated and disarmed rather than resolved numerically (§9.1). The two theories differ, in short, on what counts as a derivation, and on how much of the observed world is supposed to follow from the substrate at all. Where they make the same claim they can be compared; where Reed claims more, the comparison is not that PST scores lower on a shared axis, but that PST does not put those quantities in the derivable layer in the first place.

Causal Fermion Systems (Finster, 2010+) is the closest mathematical kin, and it does meet axes (i) and (ii): the substrate is a measure on operators of finite rank, and spacetime together with its causal structure emerges from operator-support overlap. On axis (iii) it derives a Standard-Model-*like* gauge structure, chiral fermions and weak isospin, from the causal action principle rather than the Standard Model gauge group itself, and it makes no structural commitment on the generation count or on any specific coupling value.

These five axes are, by construction, PST's own load-bearing deliverables. The table is therefore best read not as a ranking on a neutral scale, but as a

precise statement of what PST supplies that no surveyed alternative, on its own stated terms, does.

Appendix E: Bibliographic References

[1] A. Einstein, "Die Feldgleichungen der Gravitation," *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, 844 (1915). meelker.nl/pst-se/ref1

[2] R. Penrose, "Gravitational collapse and space-time singularities," *Phys. Rev. Lett.* **14** (1965) 57. doi:10.1103/PhysRevLett.14.57 · meelker.nl/pst-se/ref2

[3] S. W. Hawking and R. Penrose, "The singularities of gravitational collapse and cosmology," *Proc. R. Soc. A* **314** (1970) 529. doi:10.1098/rspa.1970.0021 · meelker.nl/pst-se/ref3

[4] L. Bombelli, J. Lee, D. Meyer and R. D. Sorkin, "Space-time as a causal set," *Phys. Rev. Lett.* **59** (1987) 521. doi:10.1103/PhysRevLett.59.521 · meelker.nl/pst-se/ref4

[5] R. D. Sorkin, "Forks in the road, on the way to quantum gravity," *Int. J. Theor. Phys.* **36** (1997) 2759. doi:10.1007/BF02435709 · meelker.nl/pst-se/ref5

[6] C. Rovelli and L. Smolin, "Loop space representation of quantum general relativity," *Nucl. Phys. B* **331** (1990) 80. doi:10.1016/0550-3213(90)90019-A · meelker.nl/pst-se/ref6

[7] A. Ashtekar, "New variables for classical and quantum gravity," *Phys. Rev. Lett.* **57** (1986) 2244. doi:10.1103/PhysRevLett.57.2244 · meelker.nl/pst-se/ref7

[8] V. L. Ginzburg and L. D. Landau, "On the theory of superconductivity," *Zh. Eksp. Teor. Fiz.* **20** (1950) 1064. Reprinted in L. D. Landau, *Collected Papers*, Pergamon Press (1965). meelker.nl/pst-se/ref8

[9] S. Mac Lane, *Categories for the Working Mathematician*, Springer, New York (1971). doi:10.1007/978-1-4612-9839-7 · meelker.nl/pst-se/ref9

[10] C. W. Misner, K. S. Thorne and J. A. Wheeler, *Gravitation*, W. H. Free-

man, San Francisco (1973). meelker.nl/pst-se/ref10

[11] H. B. G. Casimir, "On the attraction between two perfectly conducting plates," *Proc. Kon. Ned. Akad. Wetensch.* **51** (1948) 793. meelker.nl/pst-se/ref11

[12] S. K. Lamoreaux, "Demonstration of the Casimir force in the 0.6 to 6 μm range," *Phys. Rev. Lett.* **78** (1997) 5. doi:10.1103/PhysRevLett.78.5 · meelker.nl/pst-se/ref12

[13] G. Bressi, G. Carugno, R. Onofrio and G. Ruoso, "Measurement of the Casimir force between parallel metallic surfaces," *Phys. Rev. Lett.* **88** (2002) 041804. doi:10.1103/PhysRevLett.88.041804 · meelker.nl/pst-se/ref13

[14] R. S. Decca, D. López, E. Fischbach and D. E. Krause, "Measurement of the Casimir force between dissimilar metals," *Phys. Rev. Lett.* **91** (2003) 050402. doi:10.1103/PhysRevLett.91.050402 · meelker.nl/pst-se/ref14

[15] M. Bordag, U. Mohideen and V. M. Mostepanenko, "New developments in the Casimir effect," *Phys. Rept.* **353** (2001) 1. doi:10.1016/S0370-1573(01)00015-1 · meelker.nl/pst-se/ref15

[16] N. Arkani-Hamed, S. Dimopoulos and G. Dvali, "The hierarchy problem and new dimensions at a millimeter," *Phys. Lett. B* **429** (1998) 263–272. doi:10.1016/S0370-2693(98)00466-3 · meelker.nl/pst-se/ref16

[17] E. M. Lifshitz, "The theory of molecular attractive forces between solids," *Sov. Phys. JETP* **2** (1956) 73. meelker.nl/pst-se/ref17

[18] E. P. Verlinde, "On the origin of gravity and the laws of Newton," *J. High Energy Phys.* **04** (2011) 029. doi:10.1007/JHEP04(2011)029 · meelker.nl/pst-se/ref18

[19] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Phys. Rev. Lett.* **75** (1995) 1260. doi:10.1103/PhysRevLett.75.1260 · meelker.nl/pst-se/ref19

[20] A. Einstein, "Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie," *Sitzungsber. Preuss. Akad. Wiss. Berlin* (1917) 142–152. meelker.nl/pst-se/ref20

[21] E. Noether, "Invariante Variationsprobleme," *Nachr. D. König. Gesellsch. D. Wiss. Zu Göttingen, Math-phys. Klasse*, 235 (1918). English translation: M. A. Tavel, *Transp. Theory Stat. Phys.* **1** (1971) 186. doi:10.1080/00411457108231446 · meelker.nl/pst-se/ref21

[22] D. Lovelock, "The Einstein tensor and its generalizations," *J. Math. Phys.* **12** (1971) 498–501. doi:10.1063/1.1665613 · meelker.nl/pst-se/ref22

[23] P. A. M. Dirac, *The Principles of Quantum Mechanics*, Clarendon Press, Oxford (1930); 4th ed. (1958). meelker.nl/pst-se/ref23

[24] J. von Neumann, *Mathematische Grundlagen der Quantenmechanik*, Springer, Berlin (1932). English translation: *Mathematical Foundations of Quantum Mechanics*, Princeton University Press (1955; reprint 2018). meelker.nl/pst-se/ref24

[25] R. P. Feynman and A. R. Hibbs, *Quantum Mechanics and Path Integrals*, McGraw-Hill, New York (1965). Dover emended edition (2010). meelker.nl/pst-se/ref25

[26] M. B. Green, J. H. Schwarz and E. Witten, *Superstring Theory* (2 vols.), Cambridge University Press (1987). meelker.nl/pst-se/ref26

[27] E. Witten, "String theory dynamics in various dimensions," *Nucl. Phys. B* **443** (1995) 85–126. doi:10.1016/0550-3213(95)00158-O · meelker.nl/pst-se/ref27

[28] T. Kaluza, "Zum Unitätsproblem der Physik," *Sitzungsber. Preuss. Akad. Wiss. Berlin* (1921) 966–972. meelker.nl/pst-se/ref28

[29] O. Klein, "Quantentheorie und fünfdimensionale Relativitätstheorie," *Zeitschrift für Physik* **37** (1926) 895–906. doi:10.1007/BF01397481 · meelker.nl/pst-se/ref29

[30] L. Randall and R. Sundrum, "A large mass hierarchy from a small extra dimension," *Phys. Rev. Lett.* **83** (1999) 3370–3373. doi:10.1103/PhysRevLett.83.3370 · meelker.nl/pst-se/ref30

[31] J. Ambjørn, J. Jurkiewicz and R. Loll, "Reconstructing the universe," *Phys. Rev. D* **72** (2005) 064014. doi:10.1103/PhysRevD.72.064014 · meelker.nl/pst-se/ref31

[32] G. 't Hooft, "Dimensional reduction in quantum gravity," arXiv:gr-qc/9310026 (1993). meelker.nl/pst-se/ref32

[33] J. Maldacena, "The large N limit of superconformal field theories and supergravity," *Adv. Theor. Math. Phys.* **2** (1998) 231–252. doi:10.4310/ATMP.1998.v2.n2.a1 · meelker.nl/pst-se/ref33

[34] I. Kant, *Kritik der reinen Vernunft*, Hartknoch, Riga (1781); English translation *Critique of Pure Reason*, trans. N. Kemp Smith, Macmillan, London (1929). meelker.nl/pst-se/ref34

[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 · meelker.nl/pst-se/ref35

[36] J. B. Hartle and S. W. Hawking, "Wave function of the universe," *Phys. Rev. D* **28** (1983) 2960–2975. doi:10.1103/PhysRevD.28.2960 · meelker.nl/pst-se/ref36

[37] R. Geroch, "Domain of dependence," *J. Math. Phys.* **11** (1970) 437–449. doi:10.1063/1.1665157 · meelker.nl/pst-se/ref37

[38] R. M. Wald, *General Relativity*, University of Chicago Press, Chicago (1984). meelker.nl/pst-se/ref38

[39] G. Ryle, *The Concept of Mind*, Hutchinson, London (1949); Routledge reissue (2009). doi:10.4324/9780203875858 · meelker.nl/pst-se/ref39

[40] H. Putnam, *Reason, Truth and History*, Cambridge University Press,

Cambridge (1981). doi:10.1017/CBO9780511625398 · meelker.nl/pst-se/ref40

[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. meelker.nl/pst-se/ref41

[42] A. Vilenkin, "Creation of universes from nothing," *Phys. Lett. B* **117** (1982) 25–28. doi:10.1016/0370-2693(82)90866-8 · meelker.nl/pst-se/ref42

[43] J. Goldstone, "Field theories with superconductor solutions," *Nuovo Cimento* **19** (1961) 154–164. doi:10.1007/BF02812722 · meelker.nl/pst-se/ref43

[44] P. A. M. Dirac, "The cosmological constants," *Nature* **139** (1937) 323. doi:10.1038/139323a0 · meelker.nl/pst-se/ref44

[45] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), "Observation of gravitational waves from a binary black hole merger," *Phys. Rev. Lett.* **116** (2016) 061102. doi:10.1103/PhysRevLett.116.061102 · meelker.nl/pst-se/ref45

[46] J. D. Bekenstein, "Black holes and entropy," *Phys. Rev. D* **7** (1973) 2333–2346. doi:10.1103/PhysRevD.7.2333 · meelker.nl/pst-se/ref46

[47] S. W. Hawking, "Particle creation by black holes," *Commun. Math. Phys.* **43** (1975) 199–220. doi:10.1007/BF02345020 · meelker.nl/pst-se/ref47

[48] A. Connes, *Noncommutative Geometry*, Academic Press, San Diego (1994). meelker.nl/pst-se/ref48

[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 · meelker.nl/pst-se/ref49

[50] A. H. Chamseddine, A. Connes and M. Marcolli, "Gravity and the standard model with neutrino mixing," *Adv. Theor. Math. Phys.* **11** (2007)

991–1089, arXiv:hep-th/0610241. doi:10.4310/ATMP.2007.v11.n6.a3 · meelker.nl/pst-se/ref50

[51] D. Buttazzo, G. Degrassi, P. P. Giardino, G. F. Giudice, F. Sala, A. Salvio and A. Strumia, "Investigating the near-criticality of the Higgs boson," *J. High Energy Phys.* **12** (2013) 089, arXiv:1307.3536. doi:10.1007/JHEP12(2013)089 · meelker.nl/pst-se/ref51

[52] A. H. Guth, "Inflationary universe: A possible solution to the horizon and flatness problems," *Phys. Rev. D* **23** (1981) 347–356. doi:10.1103/PhysRevD.23.347 · meelker.nl/pst-se/ref52

[53] A. D. Linde, "Chaotic inflation," *Phys. Lett. B* **129** (1983) 177–181. doi:10.1016/0370-2693(83)90837-7 · meelker.nl/pst-se/ref53

[54] 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 · meelker.nl/pst-se/ref54

[55] B. Efron and C. Stein, "The jackknife estimate of variance," *Ann. Statist.* **9** (1981) 586–596. doi:10.1214/aos/1176345462 · meelker.nl/pst-se/ref55

[56] R. O'Donnell, *Analysis of Boolean Functions*, Cambridge University Press (2014). ISBN 978-1107038325. doi:10.1017/CBO9781139814782 · meelker.nl/pst-se/ref56

[57] B. S. DeWitt, *Dynamical Theory of Groups and Fields*, Gordon and Breach, New York (1965). meelker.nl/pst-se/ref57

[58] P. B. Gilkey, "The spectral geometry of a Riemannian manifold," *J. Differential Geometry* **10** (1975) 601–618. doi:10.4310/jdg/1214433164 · meelker.nl/pst-se/ref58

[59] D. V. Vassilevich, "Heat kernel expansion: User's manual," *Phys. Rept.* **388** (2003) 279–360. doi:10.1016/j.physrep.2003.09.002 · meelker.nl/pst-se/ref59

[60] B. S. DeWitt, "Quantum theory of gravity. I. The canonical theory,"

Phys. Rev. **160** (1967) 1113–1148. doi:10.1103/PhysRev.160.1113 · meelker.nl/pst-se/ref60

[61] G. Spencer-Brown, *Laws of Form*, Allen & Unwin, London (1969). meelker.nl/pst-se/ref61

[62] 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 (1990); reprinted in *Feynman and Computation*, CRC Press (2018). doi:10.1201/9780429500459-19 · meelker.nl/pst-se/ref62

[63] M. Tegmark, "The Mathematical Universe," *Foundations of Physics* **38** (2008) 101–150. doi:10.1007/s10701-007-9186-9 · meelker.nl/pst-se/ref63

[64] 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 · meelker.nl/pst-se/ref64

[65] S. French, *The Structure of the World: Metaphysics and Representation*, Oxford University Press, Oxford (2014). doi:10.1093/acprof:oso/9780199684847.001.0001 · meelker.nl/pst-se/ref65

[66] D. Bohm, *Wholeness and the Implicate Order*, Routledge & Kegan Paul, London (1980); Routledge reissue (2002). doi:10.4324/9780203995150 · meelker.nl/pst-se/ref66

[67] G. W. F. Hegel, *Science of Logic*, trans. A. V. Miller, Allen & Unwin, London (1969); original *Wissenschaft der Logik*, Cotta, Nuremberg (1812–1816). meelker.nl/pst-se/ref67

[68] A. Friedmann, "Über die Krümmung des Raumes," *Zeitschrift für Physik* **10** (1922) 377–386. doi:10.1007/BF01332580 · meelker.nl/pst-se/ref68

[69] 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 · meelker.nl/pst-se/ref69

[70] S. Perlmutter *et al.* (Supernova Cosmology Project), "Measurements of Ω and Λ from 42 high-redshift supernovae," *Astrophys. J.* **517** (1999) 565–586. doi:10.1086/307221 · meelker.nl/pst-se/ref70

[71] N. Aghanim *et al.* (Planck Collaboration), "Planck 2018 results. VI. Cosmological parameters," *Astron. Astrophys.* **641** (2020) A6. doi:10.1051/0004-6361/201833910 · meelker.nl/pst-se/ref71

[72] H. Bondi and T. Gold, "The steady-state theory of the expanding universe," *Mon. Not. R. Astron. Soc.* **108** (1948) 252–270. doi:10.1093/mnras/108.3.252 · meelker.nl/pst-se/ref72

[73] F. Hoyle, "A new model for the expanding universe," *Mon. Not. R. Astron. Soc.* **108** (1948) 372–382. doi:10.1093/mnras/108.5.372 · meelker.nl/pst-se/ref73

[74] A. A. Penzias and R. W. Wilson, "A measurement of excess antenna temperature at 4080 Mc/s," *Astrophys. J.* **142** (1965) 419–421. doi:10.1086/148307 · meelker.nl/pst-se/ref74

[75] P. W. Higgs, "Broken symmetries and the masses of gauge bosons," *Phys. Rev. Lett.* **13** (1964) 508–509. doi:10.1103/PhysRevLett.13.508 · meelker.nl/pst-se/ref75

[76] 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 · meelker.nl/pst-se/ref76

[77] S. Weinberg, "A model of leptons," *Phys. Rev. Lett.* **19** (1967) 1264–1266. doi:10.1103/PhysRevLett.19.1264 · meelker.nl/pst-se/ref77

[78] A. Salam, "Weak and electromagnetic interactions," in *Proceedings of the 8th Nobel Symposium*, ed. N. Svartholm, Almqvist & Wiksell, Stockholm (1968), pp. 367–377 [reprinted in *Selected Papers of Abdus Salam*, World Scientific Series in 20th Century Physics (1994)]. doi:10.1142/

9789812795915_0034 · meelker.nl/pst-se/ref78

[79] S. L. Glashow, "Partial symmetries of weak interactions," *Nucl. Phys.* **22** (1961) 579–588. doi:10.1016/0029-5582(61)90469-2 · meelker.nl/pst-se/ref79

[80] 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 · meelker.nl/pst-se/ref80

[81] 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 · meelker.nl/pst-se/ref81

[82] H. D. Politzer, "Reliable perturbative results for strong interactions?" *Phys. Rev. Lett.* **30** (1973) 1346–1349. doi:10.1103/PhysRevLett.30.1346 · meelker.nl/pst-se/ref82

[83] 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 · meelker.nl/pst-se/ref83

[84] 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. doi:10.1103/PhysRev.105.1413 · meelker.nl/pst-se/ref84

[85] S. L. Adler, "Axial-vector vertex in spinor electrodynamics," *Phys. Rev.* **177** (1969) 2426–2438. doi:10.1103/PhysRev.177.2426 · meelker.nl/pst-se/ref85

[86] J. S. Bell and R. Jackiw, "A PCAC puzzle: $\pi^0 \to \gamma\gamma$ in the σ-model," *Nuovo Cim. A* **60** (1969) 47–61. doi:10.1007/BF02823296 · meelker.nl/pst-se/ref86

[87] 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 · meelker.nl/pst-se/ref87

[88] 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 · meelker.nl/pst-se/ref88

[89] S. Weinberg, "Implications of dynamical symmetry breaking," *Phys. Rev. D* **13** (1976) 974–996; addendum **19** (1979) 1277–1280. doi:10.1103/PhysRevD.13.974 · meelker.nl/pst-se/ref89

[90] L. Susskind, "Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory," *Phys. Rev. D* **20** (1979) 2619–2625. doi:10.1103/PhysRevD.20.2619 · meelker.nl/pst-se/ref90

[91] S. Coleman and E. Weinberg, "Radiative corrections as the origin of spontaneous symmetry breaking," *Phys. Rev. D* **7** (1973) 1888–1910. doi:10.1103/PhysRevD.7.1888 · meelker.nl/pst-se/ref91

[92] N. Arkani-Hamed, A. G. Cohen and H. Georgi, "Electroweak symmetry breaking from dimensional deconstruction," *Phys. Lett. B* **513** (2001) 232–240. doi:10.1016/S0370-2693(01)00741-9 · meelker.nl/pst-se/ref92

[93] R. L. Workman *et al.* (Particle Data Group), "Review of particle physics," *Prog. Theor. Exp. Phys.* **2022** (2022) 083C01. doi:10.1093/ptep/ptac097 · meelker.nl/pst-se/ref93

[94] M. J. G. Veltman, "The infrared–ultraviolet connection," *Acta Phys. Polon. B* **12** (1981) 437–457. meelker.nl/pst-se/ref94

[95] 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. doi:10.1007/978-1-4684-7571-5_9 · meelker.nl/pst-se/ref95

[96] J. Polchinski, "Renormalization and effective lagrangians," *Nucl. Phys. B*

231 (1984) 269–295. doi:10.1016/0550-3213(84)90287-6 · meelker.nl/pst-se/ref96

[97] 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. meelker.nl/pst-se/ref97

[98] U. Mosco, "Composite media and asymptotic Dirichlet forms," *J. Funct. Anal.* **123** (1994) 368–421. doi:10.1006/jfan.1994.1093 · meelker.nl/pst-se/ref98

[99] M. E. Taylor, *Partial Differential Equations I: Basic Theory*, Springer, New York (1996); 2nd ed. (2011). doi:10.1007/978-1-4419-7055-8 · meelker.nl/pst-se/ref99

[100] C. McDiarmid, "On the method of bounded differences," in *Surveys in Combinatorics 1989*, London Math. Soc. Lecture Note Ser. **141**, Cambridge Univ. Press (1989), pp. 148–188. doi:10.1017/CBO9781107359949.008 · meelker.nl/pst-se/ref100

[101] W. Feller, *An Introduction to Probability Theory and Its Applications, Vol. II*, 2nd ed., Wiley, New York (1971). meelker.nl/pst-se/ref101

[102] S. C. Brenner and L. R. Scott, *The Mathematical Theory of Finite Element Methods*, 3rd ed., Springer, New York (2008). doi:10.1007/978-0-387-75934-0 · meelker.nl/pst-se/ref102

[103] C.-G. Esseen, "A moment inequality with an application to the central limit theorem," *Skandinavisk Aktuarietidskrift* **39** (1956) 160–170. doi:10.1080/03461238.1956.10414946 · meelker.nl/pst-se/ref103

[104] I. G. Shevtsova, "On the absolute constants in the Berry-Esseen type inequalities for identically distributed summands," arXiv:1111.6554 (2011). meelker.nl/pst-se/ref104

[105] P. Jordan and E. Wigner, "Über das Paulische Äquivalenzverbot," *Z. Phys.* **47** (1928) 631–651. doi:10.1007/BF01331938 · meelker.nl/pst-

se/ref105

[106] W. H. Zurek, “Decoherence, einselection, and the quantum origins of the classical,” *Rev. Mod. Phys.* **75** (2003) 715–775, arXiv:quant-ph/0105127. doi:10.1103/RevModPhys.75.715 · meelker.nl/pst-se/ref106

[107] A. Kitaev and J. Preskill, “Topological entanglement entropy,” *Phys. Rev. Lett.* **96** (2006) 110404, arXiv:hep-th/0510092. doi:10.1103/PhysRevLett.96.110404 · meelker.nl/pst-se/ref107

[108] R. F. Streater and A. S. Wightman, *PCT, Spin and Statistics, and All That*, Benjamin, New York (1964). meelker.nl/pst-se/ref108

[109] R. Haag, *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed., Springer, Berlin (1996). doi:10.1007/978-3-642-61458-3 · meelker.nl/pst-se/ref109

[110] J. Hadamard, *Lectures on Cauchy’s Problem in Linear Partial Differential Equations*, Yale University Press, New Haven (1923; Dover reprint 1952). meelker.nl/pst-se/ref110

[111] R. Courant and D. Hilbert, *Methods of Mathematical Physics, Volume II: Partial Differential Equations*, Interscience Publishers, New York (1962). doi:10.1002/9783527617234 · meelker.nl/pst-se/ref111

[112] S. W. Hawking, “The occurrence of singularities in cosmology. III. Causality and singularities,” *Proc. R. Soc. A* **300** (1967) 187–201. doi:10.1098/rspa.1967.0164 · meelker.nl/pst-se/ref112

[113] 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 · meelker.nl/pst-se/ref113

[114] C. Bär, “The Dirac operator on space forms of positive curvature,” *J. Math. Soc. Japan* **48** (1996) 69–83. doi:10.2969/jmsj/04810069 · meelker.nl/pst-se/ref114

[115] M. A. Rieffel, "Matrix algebras converge to the sphere for quantum Gromov–Hausdorff distance," *Mem. Amer. Math. Soc.* **168** (2004) 67–91. meelker.nl/pst-se/ref115

[116] M. A. Rieffel, "Gromov–Hausdorff distance for quantum metric spaces," *Mem. Amer. Math. Soc.* **168** (2004) 1–65. meelker.nl/pst-se/ref116

[117] F. Latrémolière, "The dual-modular Gromov–Hausdorff propinquity and completeness," *J. Noncommut. Geom.* **15** (2021) 347–398; arXiv:1811.04534. meelker.nl/pst-se/ref117

[118] F. Latrémolière, "The Gromov–Hausdorff propinquity for metric spectral triples," *Adv. Math.* **404** (2022) 108393. meelker.nl/pst-se/ref118

[119] T. Bhattacharyya and S. Singla, "Sequences of operator algebras converging to odd spheres in the quantum Gromov–Hausdorff distance," arXiv:2211.04321 (2022). meelker.nl/pst-se/ref119

[120] 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 · meelker.nl/pst-se/ref120

[121] C. Furey, "Generations: three prints, in colour," *J. High Energy Phys.* **10** (2014) 046. doi:10.1007/JHEP10(2014)046 · meelker.nl/pst-se/ref121

[122] C. Furey, "Three generations, two unbroken gauge symmetries, and one eight-dimensional algebra," *Phys. Lett. B* **785** (2018) 84–89. doi:10.1016/j.physletb.2018.08.032 · meelker.nl/pst-se/ref122

[123] S. Wolfram, "A class of models with the potential to represent fundamental physics," *Complex Systems* **29** (2020) 107–536. meelker.nl/pst-se/ref123

[124] G. 't Hooft, *The Cellular Automaton Interpretation of Quantum Mechanics*, Fundamental Theories of Physics, vol. 185, Springer, Cham (2016). meelker.nl/pst-se/ref124

[125] D. Oriti, "Group field theory as the second quantization of loop quan-

tum gravity," *Class. Quantum Grav.* **33** (2016) 085005. meelker.nl/pst-se/ref125

[126] R. Penrose, "Twistor algebra," *J. Math. Phys.* **8** (1967) 345–366. doi:10.1063/1.1705200 · meelker.nl/pst-se/ref126

[127] M. Cortês and L. Smolin, "Quantum energetic causal sets," *Phys. Rev. D* **90** (2014) 044035, arXiv:1308.2206. doi:10.1103/PhysRevD.90.044035 · meelker.nl/pst-se/ref127

[128] L. Hardy, "Probability theories with dynamic causal structure: A new framework for quantum gravity," arXiv:gr-qc/0509120 (2005). meelker.nl/pst-se/ref128

[129] F. Finster and J. Kleiner, "Causal fermion systems as a candidate for a unified physical theory," *J. Phys. Conf. Ser.* **626** (2015) 012020. doi:10.1088/1742-6596/626/1/012020 · meelker.nl/pst-se/ref129

[130] T. Banks, W. Fischler, S. H. Shenker and L. Susskind, "M theory as a matrix model: A conjecture," *Phys. Rev. D* **55** (1997) 5112–5128. meelker.nl/pst-se/ref130

[131] A. D. Sakharov, "Vacuum quantum fluctuations in curved space and the theory of gravitation," *Sov. Phys. Dokl.* **12** (1968) 1040–1041 [reprinted *Gen. Rel. Grav.* **32** (2000) 365–367]. doi:10.1023/A:1001947813563 · meelker.nl/pst-se/ref131

[132] T. Padmanabhan, "Thermodynamical aspects of gravity: New insights," *Rep. Prog. Phys.* **73** (2010) 046901. meelker.nl/pst-se/ref132

[133] E. P. Verlinde, "Emergent gravity and the dark universe," *SciPost Phys.* **2** (2017) 016. meelker.nl/pst-se/ref133

[134] J. Maldacena and L. Susskind, "Cool horizons for entangled black holes," *Fortschr. Phys.* **61** (2013) 781–811. meelker.nl/pst-se/ref134

[135] M. Van Raamsdonk, "Building up spacetime with quantum entanglement," *Gen. Rel. Grav.* **42** (2010) 2323–2329. meelker.nl/pst-se/ref135

[136] B. Swingle, "Entanglement renormalization and holography," *Phys. Rev. D* **86** (2012) 065007. meelker.nl/pst-se/ref136

[137] G. Vidal, "Entanglement renormalization," *Phys. Rev. Lett.* **99** (2007) 220405. meelker.nl/pst-se/ref137

[138] G. M. Dixon, *Division Algebras: Octonions, Quaternions, Complex Numbers and the Algebraic Design of Physics*, Kluwer Academic Publishers, Dordrecht (1994). doi:10.1007/978-1-4757-2315-1 · meelker.nl/pst-se/ref138

[139] T. Dray and C. A. Manogue, "Quaternionic spin," in *Clifford Algebras and their Applications in Mathematical Physics*, R. Ablamowicz and B. Fauser (eds.), Birkhäuser, Boston (2000), pp. 21–37. meelker.nl/pst-se/ref139

[140] D. Deutsch and C. Marletto, "Constructor theory of information," *Proc. R. Soc. A* **471** (2015) 20140540. meelker.nl/pst-se/ref140

[141] B. Kriger, "The Information Substrate Theory, Volume VII: Information as the Name of Differentiatedness, and the Known Measures as Projections of One Substrate," IIIR Cosmology and Theoretical Physics (2026). doi:10.5281/zenodo.20469599 · meelker.nl/pst-se/ref141

[142] C. S. Montanari, "AI-Assisted Conceptual Development of a Pre-Geometric Cosmological Model: an Exercise in AI-Assisted Conceptual Framework Generation, Paper I: Foundations and Replication Dynamics," Fermilab Technical Note FERMILAB-TM-2923-PPD (2025). doi:10.2172/3018447 · meelker.nl/pst-se/ref142

[143] J. Reed, *The Substrate Theory of Everything: A Comprehensive Technical Thesis*, independently published (2024); ISBN 9798198320833. See also *The Substrate Theory of Everything: A Technical Paper*, ISBN 9798198317031; *The Substrate Theory of Everything: A Comprehensive Textbook*, ISBN 9798198310933; and *The Substrate: A New Theory of Everything*, ISBN 9798198228726. Released for public review on 26 June

2026 with a proposed resolution of the cosmological-constant problem. meelker.nl/pst-se/ref143

[144] M. Engliš, *Toeplitz operators on Bergman-type spaces*, Ph.D. thesis, Mathematical Institute, Czechoslovak Academy of Sciences, Prague (1991). meelker.nl/pst-se/ref144

[145] K. Zhu, *Operator Theory in Function Spaces*, 2nd ed., Mathematical Surveys and Monographs vol. 138, American Mathematical Society (2007). doi:10.1090/surv/138 · meelker.nl/pst-se/ref145

[146] G. M. Dixon, *Division Algebras; Spinors; Idempotents; The Algebraic Structure of Reality*, talk given at the 2nd Mile High Conference on Nonassociative Mathematics, Denver (2009); arXiv:1012.1304 (2010). See also [138]. meelker.nl/pst-se/ref146

[147] F. Hoyle and J. V. Narlikar, "A new theory of gravitation," *Proc. R. Soc. A* **282** (1964) 191–207; see also "Mach's principle and the creation of matter," *Proc. R. Soc. A* **273** (1963) 1–11, and J. V. Narlikar, *Introduction to Cosmology*, 3rd ed., Cambridge University Press (2002). doi:10.1098/rspa.1964.0227 · meelker.nl/pst-se/ref147

[148] R. Bott, "The stable homotopy of the classical groups," *Ann. of Math. (2)* **70** (1959) 313. doi:10.2307/1970106 · meelker.nl/pst-se/ref148

[149] M. F. Atiyah, *K-Theory*, Westview Press / Addison-Wesley Advanced Book Classics, Boulder, CO (1989). [Originally Benjamin, 1967.] meelker.nl/pst-se/ref149

[150] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed., Stochastic Modelling and Applied Probability **38**, Springer, Berlin Heidelberg (2010). doi:10.1007/978-3-642-03311-7 · meelker.nl/pst-se/ref150

[151] P. R. Halmos and L. J. Savage, "Application of the Radon–Nikodym theorem to the theory of sufficient statistics," *Annals of Mathematical Statistics* **20** (1949) 225–241. doi:10.1214/aoms/1177730032 · meelker.nl/pst-

se/ref151

[152] A. Lichnerowicz, "Spineurs harmoniques," *C. R. Acad. Sci. Paris* **257** (1963) 7. meelker.nl/pst-se/ref152

[153] H. B. Lawson and M.-L. Michelsohn, *Spin Geometry*, Princeton Mathematical Series **38**, Princeton University Press (1989). meelker.nl/pst-se/ref153

[154] P. Lounesto, *Clifford Algebras and Spinors*, 2nd ed., London Mathematical Society Lecture Note Series **286**, Cambridge University Press (2001). doi:10.1017/CBO9780511526022 · meelker.nl/pst-se/ref154

[155] C. N. Yang and R. L. Mills, "Conservation of isotopic spin and isotopic gauge invariance," *Phys. Rev.* **96** (1954) 191. doi:10.1103/PhysRev.96.191 · meelker.nl/pst-se/ref155

[156] R. Bott, "Nondegenerate critical manifolds," *Ann. of Math. (2)* **60** (1954) 248. doi:10.2307/1969631 · meelker.nl/pst-se/ref156

[157] D. Hilbert, "Die Grundlagen der Physik (Erste Mitteilung)," *Nachr. Ges. Wiss. Göttingen, Math.-Phys. Kl.* 395 (1915). meelker.nl/pst-se/ref157

[158] I. Schur, "Neue Begründung der Theorie der Gruppencharaktere," *Sitzungsber. Königl. Preuss. Akad. Wiss. Berlin* 406 (1905). meelker.nl/pst-se/ref158

[159] S. L. Sobolev, "Sur un théorème d'analyse fonctionnelle," *Mat. Sbornik* **4(46)** (1938) 471. meelker.nl/pst-se/ref159

[160] H. Yukawa, "On the interaction of elementary particles. I," *Proc. Phys.-Math. Soc. Japan (3rd Ser.)* **17** (1935) 48. doi:10.11429/ppmsj1919.17.0_48 · meelker.nl/pst-se/ref160

[161] G. C. Wick, "Properties of Bethe-Salpeter wave functions," *Phys. Rev.* **96** (1954) 1124. doi:10.1103/PhysRev.96.1124 · meelker.nl/pst-se/ref161

[162] H. Weyl, "Elektron und Gravitation. I," *Z. Phys.* **56** (1929) 330.

doi:10.1007/BF01339504 · meelker.nl/pst-se/ref162

[163] E. C. G. Stueckelberg, "Die Wechselwirkungskräfte in der Elektrodynamik und in der Feldtheorie der Kernkräfte," *Helv. Phys. Acta* **11** (1938) 225. doi:10.5169/seals-110852 · meelker.nl/pst-se/ref163

[164] A. Hurwitz, "Über die Komposition der quadratischen Formen," *Math. Ann.* **88** (1923) 1. doi:10.1007/BF01448439 · meelker.nl/pst-se/ref164

[165] J. L. Walsh, "A closed set of normal orthogonal functions," *Amer. J. Math.* **45** (1923) 5. doi:10.2307/2387224 · meelker.nl/pst-se/ref165

[166] S.-T. Yau, "On the Ricci curvature of a compact Kähler manifold and the complex Monge-Ampère equation. I," *Comm. Pure Appl. Math.* **31** (1978) 339. doi:10.1002/cpa.3160310304 · meelker.nl/pst-se/ref166

[167] M. Gell-Mann, "Symmetries of baryons and mesons," *Phys. Rev.* **125** (1962) 1067. doi:10.1103/PhysRev.125.1067 · meelker.nl/pst-se/ref167

[168] P. Billingsley, *Probability and Measure*, 3rd ed., Wiley, New York (1995). meelker.nl/pst-se/ref168

[169] I. Chavel, *Eigenvalues in Riemannian Geometry*, Pure and Applied Mathematics **115**, Academic Press, Orlando (1984). meelker.nl/pst-se/ref169

[170] I. M. Gelfand and S. V. Fomin, *Calculus of Variations*, Prentice-Hall, Englewood Cliffs, NJ (1963). meelker.nl/pst-se/ref170

[171] H. Georgi, *Lie Algebras in Particle Physics*, 2nd ed., Frontiers in Physics **54**, Westview Press, Boulder, CO (1999). meelker.nl/pst-se/ref171

[172] J. W. Milnor and J. D. Stasheff, *Characteristic Classes*, Annals of Mathematics Studies **76**, Princeton University Press (1974). meelker.nl/pst-se/ref172

[173] M. Reed and B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, rev. ed., Academic Press, New York (1980).

meelker.nl/pst-se/ref173

[174] I. Newton, *Philosophiæ Naturalis Principia Mathematica*, Royal Society, London (1687). meelker.nl/pst-se/ref174

[175] H. Minkowski, "Raum und Zeit," *Phys. Z.* **10** (1909) 104. meelker.nl/pst-se/ref175

[176] V. A. Fock, "Konfigurationsraum und zweite Quantelung," *Z. Phys.* **75** (1932) 622. doi:10.1007/BF01344458 · meelker.nl/pst-se/ref176

[177] E. Majorana, "Teoria simmetrica dell'elettrone e del positrone," *Nuovo Cimento* **14** (1937) 171. doi:10.1007/BF02961314 · meelker.nl/pst-se/ref177

[178] M. F. Atiyah and I. M. Singer, "The index of elliptic operators on compact manifolds," *Bull. Amer. Math. Soc.* **69** (1963) 422. doi:10.1090/S0002-9904-1963-10957-X · meelker.nl/pst-se/ref178

[179] H. P. Robertson, "Kinematics and world-structure," *Astrophys. J.* **82** (1935) 284. doi:10.1086/143681 · meelker.nl/pst-se/ref179

[180] A. G. Walker, "On Milne's theory of world-structure," *Proc. London Math. Soc. (2)* **42** (1937) 90. doi:10.1112/plms/s2-42.1.90 · meelker.nl/pst-se/ref180

[181] H. A. Lorentz, "Electromagnetic phenomena in a system moving with any velocity smaller than that of light," *Proc. Roy. Neth. Acad. Arts Sci.* **6** (1904) 809. meelker.nl/pst-se/ref181

[182] B. Riemann, "Über die Hypothesen, welche der Geometrie zu Grunde liegen," *Abh. Königl. Ges. Wiss. Göttingen* **13** (1868) 133. meelker.nl/pst-se/ref182

[183] L. Boltzmann, "Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung," *Wiener Berichte* **76** (1877) 373. English translation: K. Sharp and F. Matschinsky, *Entropy* **17** (2015) 1971. doi:10.3390/e17041971 ·

meelker.nl/pst-se/ref183

[184] C. Hermite, "Remarque sur un théorème de M. Cauchy," *C. R. Acad. Sci. Paris* **41** (1855) 181. meelker.nl/pst-se/ref184

[185] Euclid, *The Thirteen Books of Euclid's Elements*, transl. T. L. Heath, 2nd ed., Cambridge University Press (1926). meelker.nl/pst-se/ref185

[186] B. Taylor, *Methodus Incrementorum Directa et Inversa*, Gulielmus Innys, London (1715). meelker.nl/pst-se/ref186

[187] G. W. Leibniz, "Nova methodus pro maximis et minimis," *Acta Eruditorum* 467 (1684). meelker.nl/pst-se/ref187

[188] A. Tonomura, J. Endo, T. Matsuda, T. Kawasaki and H. Ezawa, "Demonstration of single-electron buildup of an interference pattern," *Am. J. Phys.* **57** (1989) 117–120. doi:10.1119/1.16104 · meelker.nl/pst-se/ref188

[189] J. A. Wheeler, "The 'past' and the 'delayed-choice' double-slit experiment," in A. R. Marlow (ed.), *Mathematical Foundations of Quantum Theory*, Academic Press, New York (1978) 9–48. doi:10.1016/B978-0-12-473250-6.50006-6 · meelker.nl/pst-se/ref189

[190] M. O. Scully and K. Drühl, "Quantum eraser: a proposed photon correlation experiment concerning observation and 'delayed choice' in quantum mechanics," *Phys. Rev. A* **25** (1982) 2208–2213. doi:10.1103/PhysRevA.25.2208 · meelker.nl/pst-se/ref190

[191] Y.-H. Kim, R. Yu, S. P. Kulik, Y. Shih and M. O. Scully, "Delayed 'choice' quantum eraser," *Phys. Rev. Lett.* **84** (2000) 1–5, arXiv:quant-ph/9903047. doi:10.1103/PhysRevLett.84.1 · meelker.nl/pst-se/ref191

[192] W. Gerlach and O. Stern, "Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld," *Zeitschrift für Physik* **9** (1922) 349–352. doi:10.1007/BF01326983 · meelker.nl/pst-se/ref192

[193] S. Kochen and E. P. Specker, "The problem of hidden variables in quantum mechanics," *J. Math. Mech.* **17** (1967/1968) 59–87. doi:10.1512/

iumj.1968.17.17004 · meelker.nl/pst-se/ref193

[194] J. S. Bell, "On the Einstein Podolsky Rosen paradox," *Physics Physique Fizika* **1** (1964) 195–200. doi:10.1103/PhysicsPhysiqueFizika.1.195 · meelker.nl/pst-se/ref194

[195] G. C. Ghirardi, A. Rimini and T. Weber, "Unified dynamics for microscopic and macroscopic systems," *Phys. Rev. D* **34** (1986) 470–491. doi:10.1103/PhysRevD.34.470 · meelker.nl/pst-se/ref195

[196] A. Bassi, K. Lochan, S. Satin, T. P. Singh and H. Ulbricht, "Models of wave-function collapse, underlying theories, and experimental tests," *Rev. Mod. Phys.* **85** (2013) 471–527, arXiv:1204.4325. doi:10.1103/RevModPhys.85.471 · meelker.nl/pst-se/ref196

[197] Y. Y. Fein, P. Geyer, P. Zwick, F. Kiałka, S. Pedalino, M. Mayor, S. Gerlich and M. Arndt, "Quantum superposition of molecules beyond 25 kDa," *Nature Physics* **15** (2019) 1242–1245. doi:10.1038/s41567-019-0663-9 · meelker.nl/pst-se/ref197

[198] DESI Collaboration, "DESI DR2 results. II. Measurements of baryon acoustic oscillations and cosmological constraints," *Phys. Rev. D* **112** (2025) 083515, arXiv:2503.14738. doi:10.1103/tr6y-kpc6 · meelker.nl/pst-se/ref198

[199] S. Capozziello, H. Chaudhary, T. Harko and G. Mustafa, "Is dark energy dynamical in the DESI era? A critical review," *Physics of the Dark Universe* **51** (2026) 102196, arXiv:2512.10585. doi:10.1016/j.dark.2025.102196 · meelker.nl/pst-se/ref199

[200] S. Afroz and S. Mukherjee, "Hint toward an inconsistency between BAO and supernovae datasets: the evidence of redshift evolving dark energy from DESI DR2 is absent," *Phys. Rev. D* **113** (2026) 083514, arXiv:2504.16868. doi:10.1103/k59d-l795 · meelker.nl/pst-se/ref200

[201] L. Huang, R.-G. Cai and S.-J. Wang, "The DESI DR1/DR2 evidence for dynamical dark energy is biased by low-redshift supernovae," *Sci. China*

Phys. Mech. Astron. **68** (2025) 100413, arXiv:2502.04212. doi:10.1007/s11433-025-2754-5 · meelker.nl/pst-se/ref201

[202] H. Chaudhary, S. Capozziello, V. K. Sharma and G. Mustafa, "Does DESI DR2 challenge the ΛCDM paradigm?," *Astrophys. J.* **992** (2025) 194, arXiv:2507.21607. doi:10.3847/1538-4357/ae0458 · meelker.nl/pst-se/ref202

[203] D. M. Greenberger and A. Yasin, "Simultaneous wave and particle knowledge in a neutron interferometer," *Phys. Lett. A* **128** (1988) 391–394. doi:10.1016/0375-9601(88)90114-4 · meelker.nl/pst-se/ref203

[204] B.-G. Englert, "Fringe visibility and which-way information: an inequality," *Phys. Rev. Lett.* **77** (1996) 2154–2157. doi:10.1103/PhysRevLett.77.2154 · meelker.nl/pst-se/ref204

[205] S. Dürr, T. Nonn and G. Rempe, "Fringe visibility and which-way information in an atom interferometer," *Phys. Rev. Lett.* **81** (1998) 5705–5709. doi:10.1103/PhysRevLett.81.5705 · meelker.nl/pst-se/ref205

[206] M. Schlosshauer, "Decoherence, the measurement problem, and interpretations of quantum mechanics," *Rev. Mod. Phys.* **76** (2005) 1267–1305. doi:10.1103/RevModPhys.76.1267 · meelker.nl/pst-se/ref206

[207] R. M. Wald, "Black hole entropy is the Noether charge," *Phys. Rev. D* **48** (1993) R3427–R3431. doi:10.1103/PhysRevD.48.R3427 · meelker.nl/pst-se/ref207

[208] M. Planck, "Ueber das Gesetz der Energieverteilung im Normalspectrum," *Ann. Phys. (Leipzig)* **309** (1901) 553–563. doi:10.1002/andp.19013090310 · meelker.nl/pst-se/ref208

[209] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer and E. L. Wright, "The cosmic microwave background spectrum from the full COBE/FIRAS data set," *Astrophys. J.* **473** (1996) 576–587. doi:10.1086/178173 · meelker.nl/pst-se/ref209

[210] A. Einstein, "Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt," *Ann. Phys. (Leipzig)* **322** (1905) 132–148. doi:10.1002/andp.19053220607 · meelker.nl/pst-se/ref210

[211] E. O. Lawrence and J. W. Beams, "The element of time in the photoelectric effect," *Phys. Rev.* **32** (1928) 478–485. doi:10.1103/PhysRev.32.478 · meelker.nl/pst-se/ref211

[212] R. S. Decca, D. López, E. Fischbach, G. L. Klimchitskaya, D. E. Krause and V. M. Mostepanenko, "Tests of new physics from precise measurements of the Casimir pressure between two gold-coated plates," *Phys. Rev. D* **75** (2007) 077101. doi:10.1103/PhysRevD.75.077101 · meelker.nl/pst-se/ref212

[213] J. Dimock, "Dirac quantum fields on a manifold," *Trans. Amer. Math. Soc.* **269** (1982) 133–147. meelker.nl/pst-se/ref213

[214] C. Bär, N. Ginoux and F. Pfäffle, *Wave Equations on Lorentzian Manifolds and Quantization*, ESI Lectures in Mathematics and Physics (European Mathematical Society, 2007). doi:10.4171/037 · meelker.nl/pst-se/ref214

[215] H. A. Bethe and E. E. Salpeter, *Quantum Mechanics of One- and Two-Electron Atoms* (Springer, Berlin, 1957). doi:10.1007/978-3-662-12869-5 · meelker.nl/pst-se/ref215

[216] S. Navas *et al.* (Particle Data Group), "Review of particle physics," *Phys. Rev. D* **110** (2024) 030001. doi:10.1103/PhysRevD.110.030001 · meelker.nl/pst-se/ref216

[217] ATLAS Collaboration, "Measurement of inclusive jet and dijet cross-sections in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *J. High Energy Phys.* **05** (2018) 195. doi:10.1007/JHEP05(2018)195 · meelker.nl/pst-se/ref217

[218] ATLAS Collaboration, "Search for new non-resonant phenomena in high-mass dilepton final states with the ATLAS detector," *J. High Energy*

Phys. **11** (2020) 005. doi:10.1007/JHEP11(2020)005 · meelker.nl/pst-se/ref218

[219] ATLAS Collaboration, "A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery," *Nature* **607** (2022) 52. doi:10.1038/s41586-022-04893-w · meelker.nl/pst-se/ref219

[220] ATLAS Collaboration, "Constraints on the Higgs boson self-coupling from single- and double-Higgs production with the ATLAS detector using pp collisions at $\sqrt{s} = 13$ TeV," *Phys. Lett. B* **843** (2023) 137745. doi:10.1016/j.physletb.2023.137745 · meelker.nl/pst-se/ref220

[221] ATLAS Collaboration, "Constraints on Higgs boson production with large transverse momentum using $H \to b\bar{b}$ decays in the ATLAS detector," *Phys. Rev. D* **105** (2022) 092003. doi:10.1103/PhysRevD.105.092003 · meelker.nl/pst-se/ref221

[222] ATLAS Collaboration, "Measurement and interpretation of same-sign W boson pair production in association with two jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," *J. High Energy Phys.* **04** (2024) 026. doi:10.1007/JHEP04(2024)026 · meelker.nl/pst-se/ref222

[223] D. B. Malament, "The class of continuous timelike curves determines the topology of spacetime," *J. Math. Phys.* **18** (1977) 1399–1404. doi:10.1063/1.523436 · meelker.nl/pst-se/ref223

[224] D. J. Kleitman and B. L. Rothschild, "Asymptotic enumeration of partial orders on a finite set," *Trans. Amer. Math. Soc.* **205** (1975) 205–220. doi:10.1090/S0002-9947-1975-0369090-9 · meelker.nl/pst-se/ref224

[225] L. Bombelli, J. Henson and R. D. Sorkin, "Discreteness without symmetry breaking: a theorem," *Mod. Phys. Lett. A* **24** (2009) 2579–2587. doi:10.1142/S0217732309031958 · meelker.nl/pst-se/ref225

[226] S. Johnston, "Feynman propagator for a free scalar field on a causal set," *Phys. Rev. Lett.* **103** (2009) 180401. doi:10.1103/PhysRevLett.103.180401 · meelker.nl/pst-se/ref226

[227] M. Ahmed, S. Dodelson, P. B. Greene and R. D. Sorkin, "Everpresent Λ," *Phys. Rev. D* **69** (2004) 103523. doi:10.1103/PhysRevD.69.103523 · meelker.nl/pst-se/ref227

[228] N. Zwane, N. Afshordi and R. D. Sorkin, "Cosmological tests of everpresent Λ," *Class. Quantum Grav.* **35** (2018) 194002. doi:10.1088/1361-6382/aadc36 · meelker.nl/pst-se/ref228

[229] A. Eichhorn and S. Mizera, "Spectral dimension in causal set quantum gravity," *Class. Quantum Grav.* **31** (2014) 125007. doi:10.1088/0264-9381/31/12/125007 · meelker.nl/pst-se/ref229

[230] A. Eichhorn, S. Surya and F. Versteegen, "Spectral dimension on spatial hypersurfaces in causal set quantum gravity," *Class. Quantum Grav.* **36** (2019) 235013. doi:10.1088/1361-6382/ab47cd · meelker.nl/pst-se/ref230

[231] S. Carlip, "Dimension and dimensional reduction in quantum gravity," *Class. Quantum Grav.* **34** (2017) 193001. doi:10.1088/1361-6382/aa8535 · meelker.nl/pst-se/ref231

[232] S. Surya, "The causal set approach to quantum gravity," *Living Rev. Relativ.* **22** (2019) 5. doi:10.1007/s41114-019-0023-1 · meelker.nl/pst-se/ref232

[233] F. Takahashi *et al.* (Particle Data Group), "Review of particle physics," *Int. J. Mod. Phys. A* **41** (2026) 2630011. doi:10.1142/S0217751X26300115 · meelker.nl/pst-se/ref233

[234] ATLAS Collaboration, "Cross-section measurements of boosted Higgs boson production in final states with pairs of hadronically decaying τ-leptons with the ATLAS experiment," submitted to *J. High Energy Phys.* (2026), arXiv:2607.22193. meelker.nl/pst-se/ref234