foundations

MDPI

Article

Causal Quantum Spacetime Theory (CQST): A Framework for Emergent Geometry and Its Observational Test

Franco Cotana [1,2]

[1] Centro Interuniversitario di Ricerca sull'Inquinamento e sull'Ambiente (CIRIAF), Università degli Studi di Perugia, 06125 Perugia, Italy; franco.cotana@rse-web.it
[2] RSE—Ricerca sul Sistema Energetico SpA, 20134 Milano, Italy

Abstract

The Causal Quantum Spacetime Theory (CQST) treats spacetime as emergent from a discrete causal substrate, the quantum spacetime regenerating (QSR) lattice, assembling its geometric, holographic, quantum-foundational and phenomenological strands. The framework identifies the QSR lattice with Fuller's Isotropic Vector Matrix and adopts holographic scaling together with the Bekenstein–Hawking one-quarter factor. The vacuum causal density, the squared ratio of the Planck length to the de Sitter radius fixed by Λ, is a constant of spacetime; its logarithmic relation to the fine-structure constant α is a one-parameter calibration predicting no drift. The reduced Planck constant is reinterpreted as the action of one causal tick. The empirical content lies in the photon sector: the discrete electromagnetic action on the face-centred cubic lattice is built over its twelve nearest-neighbour vectors and the propagator obtained within the Brillouin zone. The linear coefficient vanishes identically and helicity-odd birefringence is excluded by inversion. In the twelve-fold FCC formulation, the leading quadratic correction is subluminal in every direction, with isotropic part 1/20 in units of the squared nearest-neighbour separation and an octahedral modulation of 5/18. The quadratic sector lies far below current sensitivity; the first two predictions are exact falsification tests of the assumed lattice symmetries.

Keywords: emergent spacetime; causal set; quantum-gravity phenomenology; Isotropic Vector Matrix; holography; causal density; fine-structure constant; Planck constant; Lorentz-Invariance violation; gamma-ray bursts

1. Introduction

Reconciling general relativity with quantum mechanics remains one of the central open problems of fundamental physics. The Causal Quantum Spacetime Theory (CQST) proposes that spacetime emerges as a statistical description of a discrete causal graph G, rather than a pre-existing continuum. Its elementary degrees of freedom are quantum spacetime regenerating (QSR) nodes, interpreted as minimal regenerating units of spacetime. Rather than representing points embedded in a manifold, QSR nodes generate large-scale geometry through connectivity, causal density and dynamical transport. Directed edges U_{ij} encode elementary causal transfer processes, while metric properties arise only after coarse-graining. In this sense CQST belongs to the wider family of approaches that derive geometry from combinatorial or causal structure, including causal-set theory, loop quantum gravity and Causal Dynamical Triangulations (CDT), while placing particular emphasis on regeneration, torsion and vacuum fixed-point dynamics [1–6].

Check for updates

Academic Editor: Ignazio Licata

Received: 30 July 2026
Revised: 24 September 2026
Accepted: 27 October 2026
Published: 3 November 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Why does the reduced Planck constant $\hbar$ have the value and dimensions it does? In standard formulations $\hbar$ is measured and inserted as an empirical constant; its consequences are enormous—stable atoms, quantized spectra, uncertainty relations, wave dynamics—yet its origin is not derived. CQST proposes that $\hbar$ is the quantitative signature of discrete causal structure at the QSR scale [2,7,8]. This should be intended (in the context of the CQST framework) as a reinterpretation rather than a derivation: the identification fixes the meaning and dimensions of $\hbar$, not its numerical value, which still enters through the lattice parameters.

The aim of this paper is to present the geometric and quantum-foundational strands as one consistent framework. Where a relation is a fit rather than a prediction, this is stated explicitly; where a proposed identification is not supported by the dynamics or observations, its limitations are reported directly.

The paper is organised as follows. Section 2 states the formal apparatus of the framework: the effective Lagrangian and the emergent metric, the discrete tick dynamics, the identification of the QSR lattice with the Isotropic Vector Matrix (IVM), the definitions of the causal density and of the fine-structure constant, the operational definition of the causal tick, and the deformed dispersion relation that supplies the observational test. Section 3 collects the consequences that follow from this apparatus, from the vacuum fixed point and holographic scaling to the C_0–α relation, the recovery of the standard quantum relations, and the confrontation of the dispersion prediction with gamma-ray-burst data. Section 4 discusses the logical structure of the framework, its speculative sectors, its limitations and the research priorities it generates. Section 5 concludes. Appendix A presents the representation-theoretic analysis of the twelve-channel bond space and its implications for lattice uniqueness.

2. Theoretical Framework

CQST is a theoretical framework rather than an experimental study; this section therefore states the postulates, definitions and mathematical structures from which every result reported in Section 3 is obtained, together with the observational quantities used to test them. No new data are generated in this work. The empirical inputs are the measured values of the Hubble constant and of the fine-structure constant, introduced in Section 2.7, and the published gamma-ray-burst and atmospheric-Cherenkov timing limits collected in Section 3.15.

2.1. The CQST Lagrangian and Emergent Geometry

In the continuum description, the QSR order parameter is represented by a real scalar field ρ, while torsion is treated as an auxiliary Einstein–Cartan degree of freedom. No curvature term is included, and no gravitational field equations are derived in this paper; the passage from the discrete dynamics to the Einstein equations is the principal missing element of the framework and is listed in Section 4.6. The corresponding effective Lagrangian is

$$\mathcal{L}_{\mathrm{CQST}} = \frac{1}{2}\,\partial_\mu\rho\,\partial^\mu\rho - V(\rho) - \frac{M_T^2}{4}\,T_{\sigma\mu\nu}T^{\sigma\mu\nu} + g_T\,\rho^2\,T_{\sigma\mu\nu}T^{\sigma\mu\nu}, \tag{1}$$

where M_T denotes the characteristic mass scale of the torsional sector. In Einstein–Cartan–Sciama–Kibble (ECSK) gravity the corresponding scale is naturally expected to be of order M_P, although no derivation is attempted within CQST. The normalisation adopted in Equation (1) makes the coupling g_T dimensionless, and the stability condition is $g_T\rho_*^2 < M_T^2/4$. The torsion field is treated throughout as an auxiliary, non-propagating degree of freedom; therefore, the quadratic term proportional to M_T^2 should be interpreted as an algebraic mass term rather than as a kinetic term.

https://doi.org/10.3390/foundations6040044

The vacuum structure of the scalar sector is determined by the Landau–Ginzburg potential

$$V(\rho) = \frac{k_2}{2}\rho^2 - \frac{k_4}{4}\rho^4 + \frac{k_6}{6}\rho^6, \tag{2}$$

and torsion is defined by

$$T^{\sigma}_{\mu\nu} = \Gamma^{\sigma}_{\mu\nu} - \Gamma^{\sigma}_{\nu\mu}, \tag{3}$$

$$\sigma^2(x,y) \equiv -2\,\ell^2_{\mathrm{QSR}} \ln \mathbb{C}(x,y), \; g_{\mu\nu}(x) = -\frac{1}{2}\,\frac{\partial^2 \sigma^2(x,y)}{\partial x^{\mu}\,\partial y^{\nu}}|_{y=x} \tag{4}$$

where $\mathbb{C}(x,y)$ denotes a connected correlator of the causal transport, normalised to unity at coincidence, $\mathbb{C}(x,x) = 1$, and decaying on the scale ℓ_{QSR}. The disconnected quantity $Tr[U(x,y)U(y,x)]$ cannot be used: for unitary transport one has $U(y,x) = U(x,y)^{-1}$ and therefore $U(x,y)U(y,x) = I$ identically, implying that its mixed second derivative vanishes. With the definition above, $\sigma^2(x,y)$ plays the role of the squared geodesic interval and the coincidence limit $y \to x$ reproduces the metric in the standard Synge form. The construction is stated here as a schematic ansatz: no proof is offered that the microscopic correlator has the assumed short-distance behaviour.

The operator content of Equations (1) and (2) is constrained by four assumed symmetries: lattice translations, which become spacetime translations in the continuum; the octahedral point group O_h, including spatial inversion; the causal orientation of the tick; and an internal $\mathbb{Z}_2$ acting as $\rho \to -\rho$, which the even-power structure of $V(\rho)$ presupposes. Within a truncation at mass dimension six, Table 1 classifies the effective operators according to their mass dimension and status, identifying those included in the present formulation and those excluded or omitted under the assumed CQST symmetries.

Table 1. Classification of the effective operators consistent with the CQST symmetry assumptions.

Operator	Dimension(s)	Status
$\partial_\mu\rho\,\partial^\mu\rho$	4	Included: kinetic term.
$\rho^2,\ \rho^4,\ \rho^6$	2, 4, 6	Included: Landau–Ginzburg potential.
$T_{\sigma\mu\nu}T^{\sigma\mu\nu}$	2	Included with the coefficient M_T^2 of Equation (1).
$\rho^2 T_{\sigma\mu\nu}T^{\sigma\mu\nu}$	4	Included: lowest-dimension $\mathbb{Z}_2$-even scalar–torsion coupling.
$\partial_\mu\rho T^\mu$, with T^μ the torsion trace	3	Excluded by the assumed $\mathbb{Z}_2$ symmetry. Without that symmetry it would be the leading scalar–torsion coupling, of lower dimension than the one retained.
$\rho^3,\ \rho^5$	3, 5	Excluded by $\mathbb{Z}_2$.
$\epsilon^{\mu\nu\lambda\sigma}T_{\mu\nu\lambda}\partial_\sigma(\rho^2)$	4	Excluded by parity, $\mathcal{P} \in O_h$. This $\mathbb{Z}_2$-even pseudoscalar coupling is forbidden by the same inversion symmetry that excludes parity-odd operators in Section 3.17, although it does not itself contain the photon field.
$\rho^2\partial_\mu\rho\partial^\mu\rho$	6	Neglected in the present formulation: it renormalises the kinetic term about the vacuum by a factor $1 + c\rho_*^2$ and can be absorbed into the normalisation of ρ at the fixed point.
$\left(T_{\sigma\mu\nu}T^{\sigma\mu\nu}\right)^2$	4	Neglected in the present formulation: a contact quartic that would modify the elimination of the auxiliary field without altering its algebraic character.
$R,\ \rho^2 R$ and all curvature operators	—	Not constructed in the present formulation: the gravitational sector is not constructed (Sections 2.1 and 4.6)

The scalar–torsion coupling retained in Equation (1) is therefore the lowest-dimension interaction admissible given the assumed $\mathbb{Z}_2$ and parity symmetries. A systematic power-counting analysis of the neglected dimension-six operators has not been carried out.

For the symmetry-breaking potential $(\rho) = \frac{k_2}{2}\rho^2 - \frac{k_4}{4}\rho^4 + \frac{k_6}{6}\rho^6$, defining $x = \rho^2$, stationary points satisfy $k_6 x^2 - k_4 x + k_2 = 0$. Real non-trivial stationary solutions therefore exist only if $k_4^2 \geq 4k_2k_6$. The non-trivial vacuum is energetically preferred over the symmetric configuration $V(0) = 0$ when $k_4^2 > \frac{16}{3}k_2k_6$, while equality corresponds to degenerate vacua. The CQST vacuum fixed point Ω_0 is assumed to lie in this broken-symmetry regime. Stability of the torsional sector additionally requires $g_T\rho_*^2 < M_T^2/4$, ensuring that the effective torsion mass-squared remains positive. The vacuum expectation value ρ_* defines the equilibrium configuration around which the low-energy effective description is constructed. Small fluctuations about this vacuum are governed by the effective parameters appearing in Equations (1) and (2), while the positivity condition above excludes tachyonic instabilities in the torsional degree of freedom.

The point of this effective description is not to assume a background metric from the start, but to treat geometry as a secondary object generated by correlations of causal transport.

Greek indices $(\mu, \nu, \sigma, \ldots)$ are reserved exclusively for spacetime tensor components. The coefficients of the Landau–Ginzburg potential are denoted by (k_2, k_4, k_6), while the metabolic relaxation rate appearing in Equation (10) is denoted by γ_{MET}. This convention prevents any collision between coupling parameters and spacetime indices. To avoid ambiguity, the torsional sector is formulated in terms of the standard Einstein–Cartan torsion tensor $T^{\lambda}_{\mu\nu}$, antisymmetric in the indices μ and ν, $T^{\lambda}_{\mu\nu} = -T^{\lambda}_{\nu\mu}$. The quadratic invariant entering Equation (1) is therefore understood as the fully contracted scalar $T_{\lambda\mu\nu}T^{\lambda\mu\nu}$, rather than as a contraction of an independent rank-two antisymmetric field. Furthermore, the symbol E is reserved exclusively for energy-related quantities $(E, E_P, \Delta E)$, while the set of graph links is denoted by L. Accordingly, causal graphs are written as $G = (V, L)$, where V is the set of QSR nodes (vertices) and L the set of directed causal links. Lastly, throughout this work, the symbol ρ is reserved for the Landau–Ginzburg scalar order parameter of the continuum theory and is not used as a tensor index. The discrete node-density variable appearing in Equation (10) is denoted by n_i and should not be identified with the scalar field ρ.

Differential geometry therefore emerges only in the continuum limit of the discrete substrate. The field ρ describes how intensely the substrate regenerates causal structure, whereas torsion captures orientational and transport asymmetries that survive coarse-graining. The effective metric is a collective correlator, not a primitive kinematic ingredient [9,10]. Three points of status should be made explicit here. First, the Lagrangian of Equations (1)–(3) is not derived from the discrete dynamics of Section 2.3. It is an effective ansatz whose admissible operator content is constrained by the symmetries of the lattice and of the causal tick: the scalar kinetic term and the sextic potential are the lowest-order terms compatible with those symmetries and with a stable non-trivial vacuum of the fixed-point structure of Section 3.1, and the scalar–torsion coupling is the lowest-dimension interaction allowed. Operators of comparable dimension have been neglected without a systematic power-counting analysis, which remains outstanding. Second, following the torsion field equations of Einstein–Cartan–Sciama–Kibble theory [9,11], the torsional sector is treated as a non-propagating auxiliary field: its equations of motion are algebraic, it carries no independent kinetic term, and enters the effective Lagrangian only through a contact contribution. The consequence for photon propagation is drawn in Section 3.17. Third, the gravitational sector is absent altogether; see Section 4.6.

2.2. The Metabolic Vacuum: Interpretive Framework

The vacuum is represented in CQST as a metabolically active state rather than an empty background. The causal graph is continuously regenerated, and the familiar continuum description of spacetime corresponds to a statistical regime in which microscopic causal events cannot be individually resolved. The QSR node is then a local unit of causal production, conservation and transfer, and the continuum manifold is a derived approximation, much as hydrodynamics is derived from molecular dynamics.

2.3. Microscopic Dynamics: The Causal Tick Operator

The effective formulation above, in terms of the scalar field ρ, the torsional sector $T^{\sigma}_{\mu\nu}$ and the emergent metric, describes the continuum limit of the theory, but it does not yet constitute the fundamental dynamics. The microscopic level is instead given by a discrete sequence of causal configurations. This sequential formulation is conceptually related to causal-set growth dynamics, in which causal structures develop through successive elementary stages [12], although the CQST update law introduced below is different in both its variables and its dynamical content:

$$G_n = (V_n, L_n, n_n, U_n), \tag{5}$$

where n counts the causal ticks. Each configuration contains the QSR nodes V_n, the oriented causal links L_n, the local causal-node density and the transport operators. The use of discrete causal updates as the fundamental dynamical substrate places CQST within the broader class of causally based approaches to quantum gravity, including energetic causal-set frameworks, although the present theory adopts a regenerating lattice structure and causal-tick dynamics distinct from those constructions [13].

The fundamental law of CQST is then the equation,

$$|\Psi_{n+1}\rangle = \hat{\Theta}_{QSR}|\Psi_{n}\rangle, \tag{6}$$

where $\hat{\Theta}_{QSR}$ is the elementary causal-tick operator. In discrete Hamiltonian form,

$$\hat{\Theta}_{QSR} = e^{(-i\hat{H}_{QSR}\tau_{QSR}/\hbar)}, \tag{7}$$

so that

$$|\Psi_{n+1}\rangle = e^{(-i\hat{H}_{QSR}\tau_{QSR}/\hbar)}|\Psi_{n}\rangle, \tag{8}$$

where τ_{QSR} denotes the duration of a single causal tick. Throughout the present formulation, τ_{QSR} is identified with the Planck time τ_P (see Section 2.8).

Here $\hbar$ enters as the fixed cost of a single causal tick (Section 2.8); the update operator therefore acts with a constant quantum of action, in keeping with the constant-$\hbar$ stance of the framework. The fundamental lattice Hamiltonian is proposed as

$$\hat{H}_{QSR} = \hat{H}_{link} + \hat{H}_{curv} + \hat{H}_{tors} + \hat{H}_{met} + \hat{H}_{reg} \tag{9}$$

The term $\hat{H}_{link}$ governs causal transport between nodes; $\hat{H}_{curv}$ measures the discrete curvature on the plaquettes; $\hat{H}_{tors}$ describes the chiral asymmetry of the cycles; $\hat{H}_{met}$ represents the local metabolic cost of the vacuum; and $\hat{H}_{reg}$ imposes regeneration and causal saturation with the constraint $K \leq 12$.

The local update law for the causal density is

$$n_i(n+1) = n_i(n) + S_i(n) - \gamma_{MET} n_i(n) + D\Sigma_{(j\in N(i))}\left[n_j(n) - n_i(n)\right] \tag{10}$$

The first term preserves the previous state, S_i is the regeneration, $\gamma_{MET} n_i$ the metabolic relaxation, D is a dimensionless diffusion coefficient governing causal transport between

neighbouring nodes, and N_i denotes the set of nearest neighbours of node i. The diffusive term therefore describes causal exchange between node i and the nodes connected to it by causal links. Writing the update in matrix form and omitting the source term, the multiplier of a mode with graph-Laplacian eigenvalue λ is $r(\lambda) = 1 - \gamma_{\mathrm{MET}} - D\lambda$. Linear stability therefore requires $0 \leq \gamma_{\mathrm{MET}} + D\lambda \leq 2$ for every eigenvalue. For the face-centred cubic graph, the Laplacian spectrum is bounded by $\lambda_{\max} = 16$, attained at the X point of the Brillouin zone, giving the stability condition $\gamma_{\mathrm{MET}} + 16D \leq 2$. Thus, declaring D dimensionless is not sufficient; D and γ_{MET} are jointly constrained by the lattice spectrum. Saturation is imposed by

$$0 \leq K_i \leq 12. \tag{11}$$

An elementary particle is then a periodic or quasi-periodic solution of the tick operator,

$$|\Psi_{n+N}\rangle = \mathrm{e}^{(-i\theta)}|\Psi_n\rangle, \tag{12}$$

and the mass emerges from the frequency of the causal cycle,

$$mc^2 = \hbar\omega = \hbar\frac{\theta}{N\tau_{\mathrm{QSR}}}. \tag{13}$$

In this form CQST finally possesses microscopic dynamics: the causal tick is the elementary operator, the discrete Hamiltonian generates its evolution, and the continuum Lagrangian of Section 2.1 is assumed to arise as the coarse-grained limit of lattice dynamics, an assumption whose present status is set out in Section 2.10. The mass relation m $c^2 = \hbar\omega$ obtained here is also the microscopic origin of the cycle-frequency mass rule (Equation (13) is written in its most general form, with θ denoting the phase accumulated over a causal cycle of N ticks. For periodic or quasi-periodic solutions, only the phase class modulo 2π is physically meaningful, so the associated quasi-energy is defined up to integer multiples of $2\pi\hbar/(N\tau_{QSR})$. Equation (13) therefore presupposes a choice of branch, taken throughout this work to be the principal branch. The familiar expression $mc^2 = \hbar(2\pi)/(N\tau_{QSR})$ is recovered for strictly periodic solutions with $\theta = 2\pi$, while quasi-periodic configurations correspond to more general values of θ).

2.4. Fuller's Synergetics and the Isotropic Vector Matrix

A major geometric claim of CQST is that the maximally symmetric QSR lattice coincides with R. Buckminster Fuller's Isotropic Vector Matrix (IVM), the discrete geometry developed in *Synergetics* [14,15]. In Fuller's construction, spheres in closest packing define an isotropic matrix when their centres are connected by equal-length vectors meeting at 60-degree angles. Removing the spheres leaves the octahedron–tetrahedron complex, the octet truss. The spatial compartments of the IVM are therefore built entirely from tetrahedra and octahedra, and the vector equilibrium is the maximally symmetric 'zero-phase' configuration from which structure can emerge.

2.5. Correspondence Between the IVM and the QSR Lattice

Within CQST the maximally symmetric QSR lattice is identified with the IVM. The status of this identification must be stated precisely: it is a defining assumption on the vacuum configuration, not a consequence of the microscopic dynamics. No theorem establishes that the axioms of Section 2.3 select the IVM, and no uniqueness result excludes other regular structures—the hyperdiamond or D_4 lattices among them—from satisfying the same conditions. The identifications below are therefore the content of the assumption, not derived correspondences. The same applies to the geometrical properties they introduce: three-dimensional closest packing, 60-degree coordination, tetrahedral–octahedral cell structure and maximal valence twelve are part of the defining assumptions of the framework. Spatial

dimensionality is likewise assumed, being built into the choice of a three-dimensional lattice, and Lorentzian signature is assumed through the causal ordering of the tick. Isotropy is not exact but approximate: as Section 3.13 shows, a residual octahedral anisotropy survives in the photon propagator, providing a geometric fingerprint of the framework, although its amplitude lies far below current observational sensitivity (Section 3.16). The framework assumes the following correspondences between the IVM and the QSR lattice:

- Closest-packed equipotential spheres in the IVM correspond to QSR nodes at approximately uniform causal density C_0;
- Equal-length geodesic vectors correspond to causal transport links U_{ij};
- The 60-degree angular coordination of the IVM corresponds to isotropic causal coordination in the graph G;
- The tetrahedral–octahedral cellular structure corresponds to the elementary plaquettes that support discrete torsion and curvature;
- Fuller's vector equilibrium corresponds to the CQST vacuum fixed point Ω_0;
- The kissing-number bound of twelve nearest neighbours corresponds to the maximal physically realisable valence of a QSR node.

2.6. The QSR Graph as an IVM Cell Complex

The maximally symmetric causal structure may be represented by a cell complex $\mathrm{K}_{\mathrm{IVM}}$ built on the graph $G = (V, L)$. The graph encodes the QSR nodes and their causal transport links, while the higher-dimensional structure of the lattice is described by tetrahedral and octahedral cells associated with closest packing. We therefore write schematically

$$K_{\mathrm{IVM}} = (V, L, C), \tag{14}$$

where $C = \Delta_{\mathrm{octahedra}} \cup \Delta_{\mathrm{tetrahedra}}$ denotes the set of tetrahedral and octahedral cells supported by the graph. The links L form the boundaries of these cells and provide the elementary causal transport structure of the lattice. In the maximally symmetric configuration, the underlying graph satisfies

$$\max[\mathrm{valence}(q_i)] = 12. \tag{15}$$

The twelve-fold coordination is attributed to the geometry of closest packing in three dimensions: if the lattice is realised through equal spheres in nearest-neighbour contact, no sphere can touch more than twelve others. Two qualifications must accompany this statement. First, the kissing-number bound of twelve holds in three dimensions, whereas in four dimensions the corresponding bound is twenty-four, attained by the D_4 lattice. The three-dimensional value is used here because the QSR lattice is a spatial graph and the fourth dimension is generated by the causal tick rather than being a direction of the lattice. This resolves the arithmetic, but at a price which is stated openly in Section 3.18: a $3 + 1$ split in which the spatial lattice is primitive introduces a preferred foliation, and hence a violation of boost invariance which the framework does not presently remove. The two questions cannot be answered independently, and Section 3.18 answers them together. Second, a preliminary analysis of two-component unitary walks on candidate lattices indicates that twelve-fold coordination may be incompatible with the unitarity conditions imposed on the tick operator, while eight-fold body-centred coordination may satisfy them (a detailed study is currently in preparation). It should be underlined that this is an unresolved tension within the framework. It is recorded here, and among the priorities of Section 4.6, rather than deferred: were it confirmed, the substrate would be body-centred rather than face-centred and the FCC calculation of Section 3.13 would have to be replaced by its BCC counterpart. The octahedral angular structure would survive, but its numerical coefficients would differ.

It is worth stating what a resolution of the uniqueness question would have to look like, since the absence of a theorem is more informative when the missing theorem is well posed.

The question is representational rather than geometrical. Let V_{12} be the twelve-dimensional space spanned by the bonds of the local cell and let Aut(C) be its automorphism group. The relevant question is whether the twelve channels constitute a single orbit of that group, and how V_{12} decomposes into irreducible representations,

$$V_{12} = \bigoplus_{\alpha} n_\alpha V_\alpha \quad \text{under Aut}(C), \quad L = D - A, \quad D = 4I_{12} \tag{16}$$

where A is the adjacency operator of the cell and L its discrete Laplacian. The twelve channels are the vertices of a cuboctahedron, with two channels adjacent when the corresponding vectors subtend an angle of sixty degrees. The resulting graph is four-regular and has twenty-four edges, which fixes the degree matrix as $D = 4I_{12}$. Appendix A provides the explicit construction and carries out the decomposition of the twelve-dimensional bond space.

The decomposition is distinctive among the cubic lattices, in the sense that the face-centred cubic (FCC), body-centred cubic (BCC) and simple cubic (SC) bond spaces decompose into different irreducible representations of O_h. This distinction does not, however, select coordination twelve: the competing lattices reproduce the same symmetry constraints on the leading isotropic and anisotropic sectors, differing only in the coefficient of the cubic harmonic. Coordination twelve therefore remains an input of the framework rather than a derived consequence. The problem nevertheless admits a definite answer, which purely geometric arguments do not, and Appendix A gives it: no representation-theoretic obstruction exists, and coordination twelve remains an assumption rather than a derived result.

The unitarity question therefore has direct consequences for the observational content of the framework. Results that follow from centrosymmetry and inversion, including the vanishing of the linear dispersion coefficient and the absence of helicity-odd birefringence, do not depend on whether the coordination is twelve-fold or eight-fold. By contrast, although the face-centred and body-centred lattices produce the same octahedral angular harmonic, its amplitude depends on the coordination. The specific angular fingerprint attributed to the IVM is therefore conditional on the admissibility of twelve-fold coordination. The quantitative consequences are derived in Sections 3.13 and 3.16.

2.7. Causal Density C_0 and the Fine-Structure Constant: Definitions

A central quantitative variable in CQST is the causal density C_0, the causal-consumption density of the cosmological vacuum,

$$C_0 = \frac{\hbar G H^2}{c^5}. \tag{17}$$

Equation (17) expresses C_0 as the squared ratio of the Planck time to a cosmological timescale. Which cosmological rate should enter is not a matter of convention and is settled below. Evaluated with the Hubble rate of the present epoch, $H_0 \approx 2.18 \times 10^{-18}$ s^{-1}, it gives

$$C_0 \sim \frac{\hbar G {H_0}^2}{c^5} \cong 1.4 \times 10^{-122}. \tag{18}$$

For order-of-magnitude purposes it is enough to note that $C_0 \sim 10^{-122}$ whichever cosmological rate is used, with the two candidates differing by a factor of order unity. The fine-structure constant is, as usual,

$$\alpha = \frac{e^2}{4\pi\varepsilon_0 \hbar c} \approx \frac{1}{137.036}. \tag{19}$$

A choice must be made here that carries more weight than its brevity suggests: which expansion rate enters the definition. The instantaneous rate H(z) is not a constant of the spacetime. It contains the matter density, which dilutes, and it therefore decreases monotonically towards an asymptotic value fixed by the cosmological constant alone. Defining

the vacuum causal density through H(z) would make a property of the vacuum depend on the amount of matter that happens to be present at a given epoch, which is difficult to motivate for a quantity intended to characterise the vacuum itself; as Section 4.5.2 shows, it also fails observationally. We therefore define C_0 through the cosmological constant,

$$H_\infty \equiv c\sqrt{\frac{\Lambda}{3}} = \sqrt{\Omega_\Lambda} H_0, \qquad C_0 \equiv (\tau_P H_\infty)^2 = \frac{\Lambda \ell_P^2}{3}, \tag{20}$$

where Λ is the cosmological constant and Ω_Λ its present density parameter. The causal density is the squared ratio of a Planck scale to a cosmological horizon scale, and may equivalently be written

$$C_0 = (\tau_P H_\infty)^2 = \left(\frac{\ell_P}{R_{dS}}\right)^2, \quad R_{dS} = \frac{c}{H_\infty} = \sqrt{\frac{3}{\Lambda}}, \tag{21}$$

with R_{dS} denoting the de Sitter horizon radius. The horizon that appears is the de Sitter radius, fixed by Λ, rather than the Hubble radius of the present epoch. The two differ by the factor $\Omega_\Lambda{}^{(-1/2)} \simeq 1.21$, so that the Hubble radius of the present epoch is a good approximation to the de Sitter radius and the physical reading of C_0—the fineness of the causal grain relative to the largest causally coherent scale—is untouched. Numerically,

$$C_0 = (\tau_P H_\infty)^2 = \Omega_\Lambda\,(\tau_P H_0)^2 \simeq 9.5 \times 10^{-123}, \quad \frac{1}{2}|\ln C_0| \simeq 140.48, \quad b \simeq 6.13 \tag{22}$$

Throughout this work, the numerical values $H_0 = 2.18 \times 10^{-18}\,\mathrm{s}^{-1}$ and $\Omega_\Lambda = 0.685$ are adopted from the Planck Collaboration 2018 cosmological parameters analysis [16]. With these parameters, the asymptotic de Sitter expansion rate is $H_\infty = \sqrt{\Omega_\Lambda}\,H_0$, and the resulting value of C_0 follows directly from its definition above.

The numerical consequences of the choice are therefore negligible, since the coefficient b of Section 3.5 is fitted in either case; what the choice settles is whether C_0 is a constant or a function of the epoch. CQST proposes that C_0 and α express the same vacuum structure, with C_0 as the natural dimensionless measure of that structure, and the definition adopted here is what allows the proposal to be stated as a relation between two constants rather than between a constant and an evolving quantity.

2.8. Operational Definition of the Causal Tick

An elementary activation of a causal edge U_{ij} is a causal tick, characterised by an energy cost $\varepsilon_{\mathrm{tick}}$ and a duration $\tau_{\mathrm{QSR}} = \ell_{\mathrm{QSR}}/c$. The reduced Planck constant is defined operationally as the action associated with a single causal tick,

$$\hbar \equiv \varepsilon_{\mathrm{tick}}\ \tau_{\mathrm{QSR}} \tag{23}$$

whose dimensions are those of action, $[\hbar]$ = energy × time. Unless otherwise stated, the fundamental QSR scales are identified with the corresponding Planck scales, $\ell_{QSR} \equiv \ell_P$, $\tau_{QSR} \equiv \tau_P$. The notation (ℓ_{QSR}, τ_{QSR}) emphasises their interpretation as microscopic lattice quantities, while (ℓ_P, τ_P) emphasises their numerical identification with the Planck length and Planck time. No distinction between the two scales is assumed in the present formulation of CQST. Action therefore acquires a direct microscopic interpretation: one causal tick corresponds to one quantum of action $\hbar$. The relation is definitional rather than predictive, since neither $\varepsilon_{\mathrm{tick}}$ nor τ_{QSR} is independently determined within the present framework. The causal tick is regarded as a discrete elementary event characterised by a duration τ_{QSR} and an associated action $\hbar$, so that Equation (23) reads $\hbar = \varepsilon_{\mathrm{tick}}\ \tau_{\mathrm{QSR}}$. The familiar quantum relations $E = \hbar\omega$ and $E = h\nu$ are recovered later in Section 3.8 for oscillatory modes composed of many causal ticks. They should therefore be interpreted as properties of periodic excitations of the lattice rather than as defining relations for an individual causal tick.

https://doi.org/10.3390/foundations6040044

2.9. Deformed Dispersion Relation: The Observational Test

The most explicit phenomenological claim is that the discrete lattice modifies photon propagation at very high energies, written as a deformed dispersion relation in powers of E/E_P, where E is the photon energy and $E_P = \sqrt{\hbar c^5/G}$ is the Planck energy:

$$E^2 = p^2c^2\left[1 + \xi\left(\frac{E}{E_P}\right) - \mathcal{A}\left(\frac{E}{E_P}\right)^2\right], \tag{24}$$

where ξ is the linear and $\mathcal{A}$ the quadratic dispersion coefficient. With this sign convention $\mathcal{A} > 0$ corresponds to subluminal propagation. The convention is fixed here and used throughout; where a signed coefficient is convenient it is written as $c_2 = -\mathcal{A}$. Different frequencies probe the lattice differently, producing energy-dependent time-of-flight effects over cosmological baselines. No phenomenological coefficients are introduced at this stage: the linear term is constrained by lattice symmetry and the quadratic coefficient is derived from the propagator calculation of Section 3.13.

2.10. From the Discrete Dynamics to the Continuum Description

The relation between the microscopic dynamics of Section 2.3 and the effective Lagrangian of Section 2.1 requires comment, since the continuum description is not derived from the discrete one in the present formulation and it would be misleading to imply otherwise. What can be offered is a coarse-graining argument valid at leading order and for the scalar sector alone.

At a schematic level, block-averaging the causal density over cells of the lattice and performing a gradient expansion of the resulting field motivates the kinetic term of Equation (1) and constrains the form of the leading potential terms. This leading-order argument is restricted to the scalar sector; it does not treat the torsional sector or determine the coefficients of the potential, which remain free parameters of the effective description. A full derivation would proceed either through a path integral over causal configurations or an explicit renormalisation procedure on the graph, and neither has been carried out.

The honest summary is therefore that the continuum Lagrangian is an effective ansatz constrained by the symmetries of the substrate, not an emergent consequence of the microscopic dynamics. This is the principal structural gap of the framework and is listed first among the theoretical priorities of Section 4.6.

2.11. Epistemic Status of the Claims of CQST

The claims assembled in this paper differ considerably in the degree of theoretical support available for each. Some are postulates adopted without derivation; some are definitions introducing a quantity without asserting anything about the world; some are mathematical consequences of the postulates; some are known relations re-expressed in the language of the framework; one is a parameter fitted to an observed value; several are conjectures; and a small number are independently falsifiable predictions. Conflating these categories is the principal risk in a framework that combines foundational assumptions, interpretive relations, mathematical consequences, and phenomenological predictions, and the status of every substantive claim is therefore stated here, before any of them is developed.

Table 2 defines the seven epistemic categories used to distinguish postulates, definitions, derived consequences, interpretations, fitted relations, ansätze, and independently falsifiable tests. In particular, the category "Interpretation" applies to relations established independently of CQST and re-expressed within its conceptual framework; without this category, such relations could be mistaken for consequences derived from CQST.

Table 2. Classification scheme for the claims of CQST.

Code	Class	Meaning
P	Postulate	A fundamental assumption. Adopted, not derived, and not derivable within the present formulation.
D	Definition	Introduces a quantity or notation. Carries no physical content by itself and cannot support a prediction.
C	Consequence	A mathematical result following from the postulates and definitions. The derivation is given in the paper.
I	Interpretation	A relation established independently of CQST, re-expressed in its terms. Constitutes no evidence for the framework.
F	Fit	A parameter chosen so as to reproduce an observed value. Agreement with that value is imposed by construction.
A	Ansatz	Motivated by symmetry or plausibility but not derived. May prove false without affecting the postulates.
T	Test	An independently falsifiable prediction, not used in fixing any parameter of the framework.

Table 3 applies this classification to the principal claims of the framework and states the basis and qualifications associated with each assignment. This classification is binding throughout the paper: no claim is presented at a level of theoretical or evidential support greater than its assigned category permits, and no quantity classified as a definition, interpretation, or fit should be regarded as evidence in favour of the framework.

Table 3. Epistemic status of the principal claims of CQST *.

Claim, Quantity or Relation	Section	Class	Basis and Qualifications
I. Substrate and microscopic dynamics			
Spacetime emergent from a discrete causal graph	Section 2.3	P	Foundational postulate of the framework
QSR lattice updated by a causal tick	Section 2.3	P	Foundational postulate of the framework
Hilbert-space state and unitary evolution of standard form	Section 2.3	P	Quantum structure is assumed at the microscopic level, not derived
Causal density C and the emergent metric	Section 2.1	D	The graph is primitive, C is defined on it, the metric is built from C
Vacuum causal density C_0 defined through Λ	Section 2.7	D	A constant of the spacetime rather than of the epoch, so α does not drift
$\hbar$ as the action of one causal tick	Section 2.8	D	Fixes the physical interpretation of $\hbar$, not its numerical value
Conformal dressing of the metric	—	—	Not used: empirically empty for any finite exponent, since $C_0 \sim 10^{-122}$
II. Effective continuum description			
Effective Lagrangian: kinetic, sextic, torsional and interaction terms	Section 2.1	A	A symmetry-constrained ansatz, not derived from the discrete dynamics
Admissible operator content fixed by lattice symmetries	Section 2.1	C	Listed in Section 2.1, with the governing symmetry for each; the Z_2 symmetry acting on ρ is assumed, not derived.

Table 3. *Cont.*

Claim, Quantity or Relation	Section	Class	Basis and Qualifications
Torsion as a non-propagating auxiliary field	Section 2.1	P	Non-propagating and algebraic, in accordance with ECSK theory
Coarse-graining of C yields the continuum scalar sector	Section 2.10	A	Leading-order coarse-graining sketch based on block-averaging and a gradient expansion; restricted to the scalar sector, with no explicit derivation or determination of the potential coefficients
Fixed point Ω_0 with positive second variation	Section 3.1	A	Assumed local stability with respect to compactly supported perturbations of the scalar sector, equipped with the H^1 norm; not demonstrated for the torsional sector or in a complete functional space.
Full continuum limit via path integral	—	—	Not obtained; the principal open problem of the framework
III. Lattice geometry			
Maximally symmetric vacuum identified with the IVM	Section 2.5	P	A defining assumption; no uniqueness theorem is available
Coordination $K = 12$; closest packing; 60° coordination	Section 2.5	P	Assumed as part of the vacuum configuration, not emergent
3 + 1 split: spatial lattice, tick generating the fourth dimension	Section 2.6	P	The three-dimensional kissing number applies to the spatial slice
Spatial dimensionality and Lorentzian signature	Section 2.5	P	Assumed; dimensionality by construction, signature via the causal tick
Exact isotropy of the vacuum	—	—	Not claimed: a residual octahedral anisotropy survives in the propagator
Admissibility of $K = 12$ under unitarity of the tick	—	—	Open; $K = 8$ may be the admissible coordination (Section 2.6)
IV. Lattice electrodynamics			
Identification of the lattice scale with the Planck length	Section 3.13	D	A provisional normalisation: fixing it before deriving the gravitational limit would be circular (Section 3.13)
Discrete electromagnetic action over the 12 nearest neighbours	Section 3.13	C	Computed in Section 3.13 over all twelve neighbours
Photon propagator from Brillouin-zone expansion	Section 3.13	C	Obtained by expansion within the Brillouin zone
Vanishing of the linear dispersion term	Section 3.16	C	Exact at every odd order. A property of the whole class of real nearest-neighbour forms on Bravais lattices, not a fingerprint of the FCC cell.
Quadratic correction carrying the octahedral harmonic	Section 3.13	C	The coefficient is direction-dependent, not a scalar
Harmonic coefficient 5/12 multiplying K_4	Section 3.13	C	Coefficient of the octahedral harmonic K_4 in the FCC angular dependence $f_{\mathrm{FCC}}(\hat{\boldsymbol{n}})$.
Subluminal propagation in every direction in the nearest-neighbour FCC formulation	Sections 3.13 and 4.6.2	C	Positive over the sphere in the current FCC formulation; conditional on the symmetry structure

Table 3. *Cont.*

Claim, Quantity or Relation	Section	Class	Basis and Qualifications
Isotropic part of the coefficient equal to 1/20	Section 3.13	C	Convention-dependent through the lattice scale; derived for the diagonal equal-weight action of Equation (62), not for general cross-bond couplings
Quadratic coefficient equal to 1/3	—	—	Not claimed: the propagator gives a value smaller by a factor of about seven
Geometric projection cot 60° as a dynamical input	—	—	Not used: cosine and sine projections are equally available
Orientation of the lattice axes in an observable frame	—	—	Not predicted; the pattern is fixed in shape but free in orientation
Absence of helicity-odd birefringence	Section 3.17	C	Parity-odd operators are excluded by the inversion element of O_h
Absence of parity-even, spatial dispersion-induced birefringence	—	—	Not claimed. Not excluded by O_h; occurs in cubic crystals at $O(k^2)$; not computed by the node-carried gauge-field formulation
Quadratic, subluminal energy-dependent photon time delay in the nearest-neighbour FCC formulation	Sections 3.16 and 4.6.2	C	Derived for the current FCC formulation; magnitude depends on lattice normalisation and sign on the symmetry structure. Beyond current observational sensitivity
Octahedral anisotropy across the sky: six axial minima, eight diagonal maxima, no dipole	Section 3.16	C	Fixed in shape but free in orientation; the amplitude presupposes coordination twelve and is conditional on the unitarity question of Section 2.6, and lies beyond current observational sensitivity
V. Lorentz symmetry			
Preferred-frame effects generated after coarse-graining	Section 3.18	C	A consequence of the 3 + 1 split, stated rather than denied
Enumeration of admissible Lorentz-violating operators	Section 3.18	C	Explicitly enumerated using the octahedral symmetry group in Section 3.18
Suppression of those operators at low energy	Section 3.18	A	Power counting in the Planck scale. Radiative stability is not established; see Section 3.18.
Recovery of exact boost invariance	—	—	Not achieved; sprinkling is not adopted, for the reasons given in Section 3.18
Radiative stability of the Planck suppression	Section 3.18	—	Not addressed: no custodial symmetry is available in the present formulation. Candidate forms are indicated in Section 4.6.2.
VI. Quantum relations			
Recovery of the Schrödinger equation	Section 3.11	I	A consistency check, not a derivation of quantum dynamics
Energy–frequency, de Broglie and uncertainty relations	Sections 3.8 and 3.11	I	These acquire a causal-tick reading; they are not derived from it
Mass and $E = mc^2$ as tick-counting identities	Sections 3.9 and 3.10	I	Reinterpretation of relations established independently of CQST
Mass as phase per tick, $m/m_p = \theta/N$	Section 3.9	C	Consequence of Equation (13) and of the closure relation; not an independent prediction.

Table 3. *Cont.*

Claim, Quantity or Relation	Section	Class	Basis and Qualifications
VII. Holography and cosmology			
Boundary-area scaling of the degrees of freedom	Section 3.2	A	Constraint counting reduces the degrees of freedom but does not fix their scaling
Bekenstein–Hawking coefficient 1/4	Section 3.3	A	An input; no counting of boundary microstates has been performed
Length scale of order the Hubble radius	Section 3.4	I	Follows algebraically from the definition of C_0
Relation between C_0 and the cosmological constant	Section 3.6	I	Follows from the standard de Sitter relation once assumed
VIII. Fine-structure relation			
Logarithmic relation between α and C_0	Section 3.5	F	A one-parameter calibration; no beta function is written down
Value $b \approx 6.13$	Section 3.5	F	Fitted to the observed α; constitutes no evidence for the framework
A value near 2π described as natural	—	—	Not claimed: no theoretical justification for any particular value is available
Derivation of b from the matter sector	—	—	Not obtained; the matter sector has not been constructed
IX. Falsifiable predictions			
Detection of a non-zero linear photon-dispersion coefficient	Section 3.16	T	CQST predicts $\xi = 0$ exactly for the centrosymmetric lattice construction; any positive detection would refute it
Detection of helicity-odd photon birefringence	Section 3.17	T	CQST predicts exactly zero by inversion symmetry; any positive detection would refute the inversion-symmetric lattice construction.
Potential sidereal modulation in terrestrial interferometry	Section 3.16	C	Consequence of a rotating terrestrial apparatus relative to the lattice frame; observability has not been established.
Coordination-independence of the isotropic coefficient 1/20 for diagonal equal-weight actions	Section 3.13	C	The coordination number cancels for the action of Equation (62); general O_h-invariant cross-bond couplings may change the result
Anisotropy amplitude as the lattice fingerprint (FCC 5/18, BCC 20/27, SC inverted)	Section 3.13	C	For the diagonal equal-weight action of Equation (62), only the anisotropy distinguishes the lattices. FCC peak-to-peak 5/18, BCC peak-to-peak 20/27, and SC peak-to-peak 10/9 with inverted ordering.
Quadratic coefficient reachable by CTA	—	—	Not claimed: the prediction lies approximately sixteen orders of magnitude below current bounds
Bound on the lattice spacing from quadratic LIV limits	Section 3.16	C	Genuine but extremely weak
Epoch dependence of α through $H(z)$	—	—	Not claimed: it would predict a drift 3–6 orders above clock, Oklo and quasar bounds (Section 4.5.2)
Absence of any drift of α with cosmic time	Section 3.5	C	Follows from the definition of C_0 through Λ; consistent with all bounds by construction

Table 3. *Cont.*

Claim, Quantity or Relation	Section	Class	Basis and Qualifications
C_0–Λ relation as an independent result	—	—	Not an independent result: it is the content of the definition of Section 2.7
Absence of helicity-odd gravitational birefringence	Section 4.5.1	C	Conditional on spin-2 tensor modes propagating on the same O_h-invariant substrate as the photon sector; under this assumption, inversion symmetry excludes the helicity-odd component
Chromatic dispersion of gravitational waves	Section 4.5.1	C	Conditional on spin-2 tensor modes propagating on the same O_h-invariant substrate; parity-even chromatic dispersion is permitted in principle, but its form has not been derived from a gravitational propagator and would be unobservable at detectable frequencies
Weak epoch dependence of α if dark energy is dynamical	Section 4.5.2	C	Conditional consequence of replacing Λ with an evolving dark-energy component in the definition of C_0; not an independent test of CQST under the constant-Λ formulation adopted in Section 2.7.
Redshift evolution of α as an experimental priority	—	—	Not a test of the framework, which predicts no evolution at all
Decomposition of the twelve-channel bond space V_{12} under Aut(C)	Appendix A	C	Computed in Appendix A: single orbit, multiplicity-free, $V_{12} \cong A_{1g} \oplus E_g \oplus T_{2g} \oplus T_{1u} \oplus T_{2u}$. The outcome is negative for uniqueness: the decomposition identifies the FCC cell but does not select it

* Classes are defined in Table 2 and section numbers refer to this paper. The final column states the basis of each entry and any qualification attaching to it. Entries marked "—" are results not obtained, or statements the framework does not make, listed here so that the open problems of the programme appear alongside its results rather than elsewhere.

3. Results

This section collects the consequences derived from the apparatus of Section 2. Throughout, results that follow from the stated postulates are distinguished from conjectures and from relations fitted to a single datum, and the distinction is made explicit at each step.

3.1. Vector Equilibrium as the Fixed Point Ω_0

The vector equilibrium of Fuller's Synergetics is interpreted as the geometric realisation of the CQST vacuum fixed point Ω_0. In the CQST framework, Ω_0 denotes the complete vacuum fixed-point configuration, whereas the vector equilibrium describes its maximally symmetric geometric manifestation within the Isotropic Vector Matrix. In both descriptions the relevant state is one of maximal symmetry, balance, and latent structure, characterised by a fixed-point map $\mathcal{F}$ ($\mathcal{F}$ denotes the dynamical fixed-point map acting on the vacuum configuration and should not be confused with the discrete plaquette curvature F_{ijk} introduced in Section 3.2). These conditions are expressed by the following fixed-point, stability, and continuum-limit relations:

$$\exists \Omega_0 \neq 0 : \mathcal{F}[\Omega_0] = \Omega_0, \tag{25}$$

$$\delta^2 S_{\text{eff}}[\Omega_0] > 0, \tag{26}$$

$$g^{(0)}_{\mu\nu} \to \eta_{\mu\nu}. \tag{27}$$

The sense in which Equation (26) expresses the assumed stability condition must be stated. Let Σ be a spatial slice of the coarse-grained description and consider perturbations $\delta\Omega \in C_c^\infty(\Sigma)$, smooth and compactly supported, equipped with the H^1 norm

$$\| \delta\Omega \|^2 = \int_\Sigma \left(| \delta\rho |^2 + \ell_0^2 | \nabla\delta\rho |^2 \right) d^3x. \tag{28}$$

Equation (26) assumes that $\delta^2 S_{\text{eff}}[\Omega_0] > 0$ for every non-zero perturbation in this class, that is, that Ω_0 is a local minimum with respect to compactly supported fluctuations of the scalar sector. No claim is made for perturbations of non-compact support, for the torsional sector, or for the existence of a global minimum (the existence and uniqueness of Ω_0 in a complete functional space, by a Banach–Schauder or direct-method argument, remains among the priorities of Section 4.6).

The vacuum is therefore not 'nothing' but a metastable organising point whose fluctuations generate observable physics. The two descriptions therefore refer to the same underlying equilibrium state viewed from complementary dynamical and geometric perspectives.

3.2. Holographic Scaling: Discrete Bianchi Constraints and the Bekenstein Bound

CQST argues that holographic behaviour need not be assumed as an external principle but arises from the acyclic and constrained nature of the causal graph. The elementary plaquettes of the lattice satisfy discrete Bianchi identities, so that bulk degrees of freedom are not all independent:

$$\sum_{\text{plaquette}} F_{ijk} = 0. \tag{29}$$

In a coarse-grained continuum limit these constraints suppress the number of freely specifiable bulk states, leaving boundary data as the independent carriers of information, so that bulk entropy satisfies a surface bound,

$$S_{\text{bulk}}(V) \leq A/4\ell_P{}^2. \tag{30}$$

Closed causal loops in the bulk would overcount information and violate the holographic bound; the absence of causal cycles is thus treated as the combinatorial origin of holographic scaling [17,18].

3.3. The One-Quarter Factor (Conjecture)

The factor one quarter in the Bekenstein–Hawking formula is conceptually striking;

$$S_{BH} = \frac{A}{4\ell_P{}^2}. \tag{31}$$

CQST offers a conjectural geometric interpretation based on the tetrahedral–octahedral structure of the IVM lattice. The elementary plaquettes of the lattice are triangular rather than quadrilateral, each being characterised by three edge variables and subject to a discrete Bianchi-type closure constraint. Such constraints reduce the number of independent bulk degrees of freedom and suggest that only a restricted subset of boundary variables contributes independently to horizon information.

In this picture, the coefficient 1/4 is not derived from a simple combinatorial counting of edges per plaquette. Rather, it is interpreted as a possible emergent consequence of the symmetry and constraint structure of the discrete lattice. The tetrahedral–octahedral geometry naturally introduces redundancy relations among neighbouring plaquettes, which may reduce the effective number of independent degrees of freedom available on a horizon surface.

We stress that this remains a heuristic geometric interpretation rather than a derivation. A rigorous counting of the independent boundary degrees of freedom of the tetrahedral–octahedral lattice is still required before the Bekenstein–Hawking coefficient can be claimed to emerge from first principles. This task is listed among the theoretical priorities in Section 4.6.

3.4. Emergent de Sitter Horizon and the Hubble Scale

Given that the observed cosmological constant is positive (Section 3.6), the relevant maximally symmetric geometry is de Sitter. Therefore, the cosmological holography is framed through the de Sitter/apparent (Gibbons–Hawking) horizon rather than through the Anti de Sitter–Conformal Field Theory (AdS/CFT). At the vacuum fixed point Ω_0 the curvature radius is governed by the causal density of the vacuum,

$$\ell_{dS} \propto \frac{\ell_P}{\sqrt{C_0}}. \tag{32}$$

Because C_0 is extremely small in Planck units, the curvature radius is of the order of the de Sitter horizon radius, and the cosmological horizon carries a Gibbons–Hawking entropy of order $A/4{\ell_p}^2$. With the definition of Section 2.7 the emergent length scale is the de Sitter radius itself, a constant of the spacetime; it exceeds the present Hubble radius by the factor $\Omega_\Lambda^{-1/2} \simeq 1.21$, so their near-coincidence is a feature of the present epoch and not an identity. This should not be presented as a prediction in either case: a length scale defined in terms of Λ and Planck quantities necessarily returns a length scale of that order, and the content of the statement is dimensional rather than dynamical. We note explicitly that de Sitter holography is far less understood than its AdS counterpart [19]; the extension of holographic arguments to this setting is therefore taken here as a working assumption.

3.5. Logarithmic Relation Between C_0 and α (One-Parameter)

CQST proposes a logarithmic relation between the electromagnetic coupling and the vacuum causal density, of the form suggested by renormalization-group (RG) running between an ultraviolet and an infrared scale. It should be said at the outset that this is a formal analogy and not an RG calculation. No beta function is written down, no field content is specified for the running, and the relation is not obtained from an action principle; the matter sector of CQST has not been constructed, so there is at present nothing to renormalize. The Standard Model, moreover, possesses multiple thresholds between the Planck and Hubble scales, so a single logarithm spanning some sixty orders of magnitude cannot be read as a running coupling in the technical sense. With those qualifications, the relation reads

$$\alpha^{-1} = \frac{b}{2\pi}\ln\left(\frac{1}{\sqrt{C_0}}\right) = \frac{b}{2\pi}\ln\left(\frac{m_P c^2}{\hbar H_\infty}\right). \tag{33}$$

The form with H_0 differs by—$1/2 \ln \Omega_\Lambda \simeq 0.19$ in a logarithm of about 140 and may be used as an approximation, provided it is labelled as such. One structural point should be recorded here. Because C_0 is defined in Section 2.7 through the cosmological constant and not through the instantaneous expansion rate, it is a constant of the spacetime, and the relation above therefore predicts no drift of α with cosmic time. This is not a restriction imposed to evade a bound but a consequence of which quantity is taken to characterise the vacuum; the alternative definition, and the observational conflict it generates, are analysed in Section 4.5.2.

This is a single-parameter relation, not a parameter-free prediction: b is presently an input tuned to one datum (α). It becomes predictive only once b is derived independently from the matter sector (Section 4.6). With that caveat, the relation retains the following qualitative features: in the flat-vacuum limit $C_0 \to 0$ the theory gives $\alpha \to 0$, while at Planckian

density $C_0 \to 1$ the coupling formally enters a strong-coupling regime. The observed α is thus presented as an infrared shadow of the causal vacuum. At the present stage, however, this statement should not be read as an independent prediction of α from C_0. Since b is fixed by the observed value of α, Equation (33) is best understood as a one-parameter recalibration of the infrared electromagnetic coupling in terms of the vacuum causal density. It would become a genuine relation between independent observables only if b were derived from the matter sector without using α as an input.

3.6. Relation to the Cosmological Constant

The relation between C_0 and the cosmological constant is not an additional result of the framework but the content of the definition adopted in Section 2.7. Since C_0 is Λ in Planck units up to a factor of three, the relation is exact by construction and carries no independent information. It is stated here because the logical order matters: Λ is primary and C_0 is defined from it, rather than the reverse. Everything the framework says about the vacuum causal density is therefore a statement about the cosmological constant expressed in Planck units. Since C_0 is defined through the cosmological constant in Section 2.7, the fitted relation of Section 3.5 may equally be expressed as a relation involving Λ:

$$\Lambda = \frac{3C_0}{{\ell_P}^2} \simeq 2.8 \times 10^{-122} {\ell_P}^{-2}. \tag{34}$$

Combined with Section 3.5, this gives

$$\alpha \sim 1/\ln\left(1/\sqrt{(\Lambda {\ell_P}^2)}\right) \sim 1/\ln\left(1/\sqrt{C_0}\right). \tag{35}$$

The intended interpretation is that the strength of electromagnetism and the cosmological constant may be parallel infrared manifestations of the same vacuum fixed-point structure. Two things must be said about the status of this. It is not an explanation of the smallness of Λ: the small number is put in through the definition of C_0 and reappears in the logarithm, and no dynamical mechanism drives it to that value. Furthermore, it remains a calibrated one-parameter correspondence rather than a parameter-free prediction, as Section 3.5 states. What the definition of Section 2.7 does secure is that the correspondence is between two constants of nature, so that it is at least well posed as a relation; under the alternative definition through $H(z)$ it was not, for the reasons set out in Section 4.5.2.

3.7. Action as Tick Counting and Natural Lattice Units

Once $\hbar$ is the cost of a causal tick, action is discrete: each elementary arc contributes one quantum of action and a finite process is characterised by the number of ticks required,

$$\delta S_{ij} = \hbar, \tag{36}$$

$$S_{\mathrm{QSR}} = \hbar \times (\text{number of causal ticks}). \tag{37}$$

Two quantities must be kept distinct here, since they differ by a factor of 2π. Equation (37) is a statement about the substrate: each activation of a causal edge contributes the action $\hbar$ defined in Equation (23), and the total is obtained by counting activations. The action of a propagating state, by contrast, is $\hbar$ times the phase it accumulates, so that a state closing a full cycle of causal phase carries action $2\pi\hbar = h$ regardless of how many ticks the closure required. The two coincide only for a state whose phase advances by one radian per tick, and Equation (37) should not be applied to a closed resonance in place of the closure condition.

The Planck units then appear as natural lattice units,

$$\ell_P = \sqrt{\frac{\hbar G}{c^3}}, \tag{38}$$

$$\tau_P = \sqrt{\frac{G\hbar}{c^5}}, \tag{39}$$

$$m_P = \sqrt{\frac{\hbar c}{G}}, \tag{40}$$

$$E_P = \sqrt{\frac{\hbar c^5}{G}}, \tag{41}$$

with closure relation

$$\hbar = E_P \tau_P = \left(m_P c^2\right)\frac{\ell_P}{c}, \tag{42}$$

so that $\hbar$ is the action of a single QSR tick, ℓ_P the size of a node, τ_P the duration of a tick, and E_P the energy of a single causal event.

3.8. $E = h\nu$ as a Consistency Relation

A photon of frequency ν is an oscillatory mode of causal transport. Each cycle is a complete sequence of ticks and the mode energy is the causal cost accumulated per unit time,

$$E = \hbar\omega = h\nu. \tag{43}$$

On this reading Planck's relation is recovered as a consequence of lattice discreteness rather than re-derived from a deeper constant: a mode of frequency ν performs ν cycles per second, and each complete cycle closes a full 2π of causal phase and therefore carries action $h = 2\pi\hbar$. Equivalently, the action accumulated per tick is $\hbar$ per radian of phase advance, so that the cost of a tick is $\hbar$ only for a mode whose phase advances by one radian per tick [7,20].

3.9. Mass as a Causal Metabolic Rate

The same reasoning applies to inertial mass. The Compton frequency is read as the rate at which causal phase advances in a QSR structure,

$$\omega_C = \frac{mc^2}{\hbar}, \tag{44}$$

$$m = \frac{\hbar\omega_C}{c^2} = \frac{\hbar}{\lambda_C c}, \tag{45}$$

where λ_C denotes the reduced Compton wavelength $\lambda_C = \hbar/(mc)$. With this convention, Equation (45) follows directly from $\omega_C = mc^2/\hbar$ and remains consistent with the standard relation $\omega_C = c/\lambda_C$. The ordinary Compton wavelength $h/(mc)$ differs by a factor 2π. For a strictly periodic solution of N ticks with total phase $\theta = 2\pi$, Equations (13) and (42) give a compact form of the same statement,

$$\frac{m}{m_P} = \frac{\theta}{N} = \Delta\varphi_{\text{tick}}, \tag{46}$$

the phase advance per tick, which for the fundamental periodic case equals $2\pi/N$. Mass in CQST is therefore phase per tick, in units of the Planck mass. This is a consequence of Equation (13) and of the closure relation of Equation (42), not an independent prediction.

Mass is thus an emergent rate—how many causal relations per unit time sustain the local QSR pattern of a particle—and the local causal energy density u_{QSR} tracks the local tick-density rate according to

$$u_{\text{QSR}}(x,t) = \hbar\dot{n}_{tick}(x,t). \tag{47}$$

Matter is not placed into spacetime from outside; it is a concentration of causal metabolism within the lattice. The framework also suggests that effective mass could

vary with local causal-density field $C(x,t)$, a speculative possibility constrained below ($\dot{n}_{tick}(x,t)$ denotes the local tick-rate density, i.e., the number of causal ticks per unit volume and unit time). Equation (47) therefore defines u_{QSR} as an energy density rather than a mass density. The corresponding mass density is obtained through the standard conversion $\rho_m = u_{QSR}/c^2$. Throughout CQST, local matter content is understood as a concentration of causal energy associated with the underlying tick structure.

3.10. $E = mc^2$

A particle at rest is still metabolically active: its node advances in causal phase at the Compton angular frequency, so the rest energy is the action per unit phase times the rate of phase advance,

$$E_0 = \hbar\omega_C = mc^2. \tag{48}$$

Mass–energy equivalence is then an algebraic identity once mass is defined by the tick rate [21]. For velocity v,

$$E_{tot} = \hbar\gamma\omega_C = \gamma mc^2, \tag{49}$$

and the Lorentz factor is the manifestation of time dilation in the bookkeeping of causal ticks.

3.11. Uncertainty, de Broglie and Schrödinger Dynamics

If $\hbar$ is the minimum cost of a causal relation, a measurement cannot cost less than one tick, giving an operational reading of the energy-time relation,

$$\Delta E \Delta t \geq \frac{\hbar}{2}, \tag{50}$$

and, from the impossibility of localization below one causal edge,

$$\Delta x \Delta p \geq \frac{\hbar}{2}. \tag{51}$$

The de Broglie relation is the causal-path length of a full phase cycle of momentum p,

$$\lambda_{\mathrm{dB}} = \frac{h}{p}, \tag{52}$$

and, in the large-number statistical limit of many QSR events, the causal flow on Hilbert space is argued to reduce to the Schrödinger equation,

$$i\hbar\partial_t|\Psi\rangle = H|\Psi\rangle, \tag{53}$$

with $|\Psi|^2$ measuring the expected density of ticks in a region [22–25]. The causal-tick interpretation provides an intuitive lower bound of order $\hbar$ by asserting that a measurement cannot involve less than one elementary causal update. The precise coefficient $\frac{1}{2}$ appearing in Equations (50) and (51), however, is inherited from the standard quantum–mechanical uncertainty relations derived from operator commutators. CQST should therefore be understood as providing a causal interpretation of the uncertainty principle rather than an independent derivation of the exact numerical prefactor.

Starting from the discrete update law in Equations (6)–(8), the continuum Schrödinger equation is recovered in the formal limit in which the causal tick interval τ_{QSR} is much smaller than the characteristic timescale of the coarse-grained state. Expanding the one-step evolution operator to first order in τ_{QSR} gives

$$\left|\Psi\left(t+\tau_{\mathrm{QSR}}\right)\right\rangle = \exp\left(\frac{-i\hat{\mathrm{H}}_{\mathrm{QSR}}\tau_{\mathrm{QSR}}}{\hbar}\right)\left|\Psi(t)\right\rangle, \tag{54}$$

then,

$$\left|\Psi\left(t+\tau_{\mathrm{QSR}}\right)\right\rangle \approx \left[1-\left(\frac{i}{\hbar}\right)\hat{\mathrm{H}}_{\mathrm{QSR}}\tau_{\mathrm{QSR}}\right]\left|\Psi(t)\right\rangle; \tag{55}$$

therefore,

$$\frac{\left|\Psi\left(t+\tau_{\mathrm{QSR}}\right)\right\rangle-\left|\Psi(t)\right\rangle}{\tau_{\mathrm{QSR}}} \approx -\left(\frac{i}{\hbar}\right)\hat{\mathrm{H}}_{\mathrm{QSR}}|\Psi(t)\rangle. \tag{56}$$

Taking the continuum limit

$$\lim_{\tau_{\mathrm{QSR}}\to 0}\frac{\left|\Psi\left(t+\tau_{\mathrm{QSR}}\right)\right\rangle-|\Psi(t)\rangle}{\tau_{\mathrm{QSR}}} = \frac{\partial}{\partial t}|\Psi(t)\rangle, \tag{57}$$

which yields

$$\frac{\partial}{\partial t}|\Psi(t)\rangle = -\left(\frac{i}{\hbar}\right)\hat{\mathrm{H}}_{\mathrm{QSR}}|\Psi(t)\rangle, \tag{58}$$

the expansion requires $\| \hat{H}_{\mathrm{QSR}}\tau_{\mathrm{QSR}}/\hbar \| \ll 1$ on the coarse-grained modes considered, or, equivalently, that the state lie in the domain of $\hat{H}_{\mathrm{QSR}}$.

This should be understood as a consistency argument: it shows how the standard Schrödinger form can emerge from the discrete tick operator under smooth-limit assumptions. It does not yet constitute proof that the microscopic QSR dynamics is necessarily linear and unitary; that remains one of the theoretical tasks of the framework.

3.12. The Speed of Light as the Propagation Speed of a Tick

CQST reinterprets the invariant speed c as the propagation speed of a single tick,

$$c = \frac{\ell_{\mathrm{QSR}}}{\tau_{\mathrm{QSR}}}. \tag{59}$$

The microscopic causal speed $c = \ell_{\mathrm{QSR}}/\tau_{\mathrm{QSR}}$ limits propagation to one elementary causal link per update. This microscopic bound does not by itself guarantee that the group velocity derived from every effective lattice dispersion relation remains below its low-energy limiting value; the latter depends on the lattice action and symmetry structure, as discussed in Section 4.6.2.

3.13. The Photon Propagator on the FCC Lattice

The angular coordination of the IVM selects a characteristic geometric ratio. Using the 60-degree coordination of the transport vectors, the projection of a transport vector onto the propagation direction scales as cot 60°, giving

$$r_{\mathrm{IVM}}(\text{geometric}) = \frac{1}{\sqrt{3}} \approx 0.577. \tag{60}$$

Equation (60) describes the geometry of the tetrahedral cell. It carries no dynamical weight and is not used to fix any observable coefficient. It does not provide a unique basis for determining a dispersion coefficient. The projections cos 60° and sin 60° are equally admissible and lead to different numerical values, while no physical principle selects the cotangent projection over the alternatives.

In any lattice-regularised field theory, the high-momentum behaviour of a propagating field is governed by the Brillouin-zone structure of the lattice, and the dispersion coefficient must be extracted from the propagator. The remainder of this section carries out that extraction.

Writing a for the side of the conventional cubic cell, the twelve nearest-neighbour vectors of the face-centred cubic lattice underlying the IVM are

$$\boldsymbol{e}_i = \frac{a}{2}(\pm 1, \pm 1, 0),\ \ \frac{a}{2}(\pm 1, 0, \pm 1),\ \ \frac{a}{2}(0, \pm 1, \pm 1),\ \ i = 1, \ldots, 12 \tag{61}$$

with nearest-neighbour separation $d = a/\sqrt{2}$. The gauge field is carried by the nodes, and the simplest lattice action reducing to the Maxwell action in the naïve continuum limit is

$$S_{\mathrm{EM}}^{\mathrm{lat}} = -\frac{1}{4}\,\frac{1}{N}\sum_{x\in\Lambda}\sum_{i=1}^{12}\left[A_\mu(x+\boldsymbol{e}_i) - A_\mu(x)\right]\left[A^\mu(x+\boldsymbol{e}_i) - A^\mu(x)\right] \tag{62}$$

where N is fixed by requiring the leading term of the dispersion relation to be c^2k^2. The construction follows the general framework of lattice field theory pioneered by Wilson, in which continuum gauge dynamics is recovered from a discrete action defined on a regular lattice through its long-wavelength behaviour and momentum-space propagator structure [26]. Since the octahedral group O_{h} of the lattice contains spatial inversion, no parity-odd contraction of the link differences is admissible; this fact is used again in Section 3.17. Upon Fourier transforming, the quadratic form is governed by the structure function

$$\sum(\boldsymbol{k}) = \sum_{i=1}^{12}[1-\cos(\boldsymbol{k}\cdot\boldsymbol{e}_i)] = 4[3-(\cos u\cos v\ +\ \cos u\cos w\ +\ \cos v\cos w)] \tag{63}$$

with $u = ak_x/2$, $v = ak_{y/2}$ and $w = ak_z/2$. The closed form on the right follows from pairing the twelve vectors into six antipodal pairs and applying the cosine addition formula; it is exact over the whole Brillouin zone, not merely at small momentum. Expanding to fourth order gives

$$\frac{\sum(\boldsymbol{k})}{a^2} = k^2 - a^2\left[\frac{1}{48}\sum_i k_i^4 + \frac{1}{16}\sum_{i<j}k_i^2k_j^2\right] + O(a^4k^6). \tag{64}$$

The quartic term contains the two independent quartic invariants of the octahedral group—the sum of fourth powers of the Cartesian components, and the sum of their pairwise products—with the distinct coefficients 1/48 and 1/16. It is precisely this inequality that prevents the correction from being isotropic: had the two coefficients stood in the ratio required by rotational invariance, the quartic term would have collapsed to a multiple of k^4 and no anisotropy would survive. Writing $\boldsymbol{k} = k\hat{\boldsymbol{n}}$ and using the unit-vector constraint to eliminate the pairwise invariant, the quartic term separates into an isotropic part and the lowest non-trivial cubic harmonic,

$$K_4(\hat{\boldsymbol{n}}) \equiv \sum_{i=1}^{3} n_i^4 - \frac{3}{5}, \quad \langle K_4\rangle_{S^2} = 0, \tag{65}$$

which is the $l = 4$ harmonic of the octahedral group and averages to zero over the sphere. The dispersion relation becomes

$$\omega^2 = c^2k^2\left[1 - a^2k^2\left(\frac{1}{40} - \frac{K_4(\hat{\boldsymbol{n}})}{96}\right)\right], \tag{66}$$

or, isolating the anisotropy as a fractional modulation of the isotropic term,

$$\omega^2 = c^2k^2\left[1 - \frac{a^2k^2}{40}\left(1 - \frac{5}{12}K_4(\hat{\boldsymbol{n}})\right)\right]. \tag{67}$$

Two features of Equation (66) deserve emphasis. The coefficient is positive for every direction on the sphere, so propagation is subluminal throughout: there is no direction along which the lattice produces superluminal photons. For the diagonal equal-weight nearest-neighbour action of Equation (62), the harmonic coefficient between the anisotropic and isotropic parts is 5/12, independently of the lattice spacing and of the normalisation N. The octahedral form of the angular dependence is fixed by symmetry, whereas the numerical coefficient 5/12 depends on the chosen lattice action and may be modified by the O_h-invariant cross-bond couplings discussed in Appendix A.5.

To express the result observationally, the lattice scale must be tied to the Planck scale, and here an ambiguity enters which should not be concealed. Identifying the nearest-neighbour separation with the Planck length, $d = \ell_p$, gives $a^2 = 2\ell_p{}^2$ and hence, with $E = \hbar ck$,

$$\omega^2 = c^2k^2\left[1 - \mathcal{A}(\hat{\boldsymbol{n}})\left(\frac{E}{E_P}\right)^2\right], \mathcal{A}(\hat{\boldsymbol{n}}) = \frac{1}{20}\left(1 - \frac{5}{12}K_4(\hat{\boldsymbol{n}})\right) \tag{68}$$

so that the coefficient takes the values

$$\mathcal{A}_{[100]} = \frac{1}{24}, \quad \mathcal{A}_{[110]} = \frac{5}{96}, \quad \mathcal{A}_{[111]} = \frac{1}{18}, \langle\mathcal{A}\rangle = \frac{1}{20}. \tag{69}$$

Had the conventional cell side rather than the nearest-neighbour separation been identified with the Planck length, every value above would be halved. The overall magnitude of the coefficient is therefore fixed by a convention relating the lattice scale to the Planck scale and is not a parameter-free consequence of the geometry; the angular structure is. The two are classified separately in Table 3.

The identification is not merely a choice of units: the primitives of the construction are an elementary causal separation ℓ_0 and a tick τ_0, and the limiting speed follows from them kinematically; the gravitational constant does not,

$$c = \frac{\ell_0}{\tau_0} \text{ (follows from the primitives)}, G \nleftrightarrow \{\ell_0, \tau_0\} \text{ (does not)} \tag{70}$$

so that setting $\ell_0 = \ell_p$ presupposes a derivation of the gravitational sector—the passage from the discrete dynamics to the Einstein equations—which the framework does not possess. Were that identification made before such a derivation, and were the gravitational limit subsequently derived using it, the argument would be circular. The identification is therefore conceptualised here as a provisional normalisation and is marked as such in Table 3, and the observational statements of Section 3.16 are framed so as not to depend on it: the bound obtained there constrains ℓ_0 rather than assuming it, and the angular pattern is independent of it altogether.

The following result establishes that for the diagonal equal-weight nearest-neighbour action of Equation (62), the isotropic coefficient is independent of the coordination number and therefore of the FCC–BCC choice discussed in Section 2.6. Consider a centrosymmetric neighbour set whose bonds all have the same length ℓ_0, with no cross-bond couplings and with the structure function normalised by its second moment as above. That normalisation gives $N = Z\ell_0{}^2/6$, where Z is the coordination number, while the spherical average of the fourth moment is $Z\ell_0{}^4k^4/5$ for every such graph. The coordination number cancels between the two, leaving

$$\langle\mathcal{A}\rangle_{S^2} = \frac{1}{24}\cdot\frac{Z\ell_0^4/5}{Z\ell_0^2/6}\cdot\frac{1}{\ell_0^2} = \frac{1}{20} \quad \text{(independent of Z)} \tag{71}$$

Within this class of diagonal equal-weight nearest-neighbour actions, the isotropic coefficient 1/20 is independent of the coordination number. It is therefore shared by the FCC ($Z = 12$), BCC ($Z = 8$), and SC ($Z = 6$) graphs, each of which returns 0.050000. This result does not extend to the general O_h-invariant cross-bond couplings discussed in Appendix A.5. What distinguishes one lattice from another is the anisotropy alone:

$$f_{\mathrm{FCC}}(\hat{\boldsymbol{n}}) = 1 - \frac{5}{12}K_4, \; f_{\mathrm{BCC}}(\hat{\boldsymbol{n}}) = \frac{5}{9}\left(\frac{9}{5} - 2K_4\right) = 1 - \frac{10}{9}K_4, \quad f_{\mathrm{SC}}(\hat{\boldsymbol{n}}) = 1 + \frac{5}{3}K_4. \tag{72}$$

The body-centred cubic graph therefore gives a peak-to-peak anisotropy of 20/27, that is, 74.1 per cent, nearly three times the face-centred value, while for the simple cubic graph the angular ordering is inverted, with retardation being strongest along the cube axes rather than weakest. The amplitude and orientation of the octahedral pattern, and not the

magnitude of the isotropic term, constitute the fingerprint of the lattice. This comparison makes the dependence of the angular fingerprint on the unitarity question explicit. For twelve-fold FCC coordination, the harmonic coefficient is 5/12, the normalised peak-to-peak modulation is 5/18, and the coefficients along the [100], [110], and [111] directions are in the ratio 12 : 15 : 16. For eight-fold BCC coordination, the corresponding values are 10/9, 20/27, and 3 : 6 : 7, respectively. The two substrates therefore produce the same octahedral harmonic and angular ordering, but with different amplitudes. The orientation of the pattern cannot distinguish between them; only its amplitude can. For the diagonal equal-weight action considered here, the isotropic result remains unchanged whichever coordination the unitarity analysis of Section 2.6 ultimately selects, whereas the lattice fingerprint depends on the FCC–BCC choice.

3.14. Observational Time Delay at Linear Order

This section records the coefficient that the centrosymmetry theorem of Section 3.16 will show to vanish identically. Two photons of different energy emitted simultaneously from a distant source would arrive with a linear-order delay

$$\Delta t = \xi \left[\left(\frac{E_1 - E_2}{E_P} \right) \right] \frac{L}{c}. \tag{73}$$

For a representative burst with $\Delta E \cong 10$ GeV and $d \cong 1$ Gpc, the estimate is

$$\Delta t \cong \xi \times 0.084 \text{ s}. \tag{74}$$

If the bare IVM geometric ratio were naively identified with a linear observable coefficient, $\xi = r_{\text{IVM}} = 1/\sqrt{3}$, the corresponding delay would be

$$\Delta t \cong 0.577 \times 0.084 \text{ s} \cong 4.9 \times 10^{-2} \text{ s}, \tag{75}$$

namely about 49 ms. This is not a separate estimate, but the direct numerical specialisation of Equation (73). Its significance is that even a sub-second effect at GeV energies can be constrained by gamma-ray-burst timing, as discussed in Section 3.15.

3.15. Confrontation with Data: Exclusion of the Geometric Linear Coefficient

Comparing with published limits in the linear scenario:

- Fermi-LAT, GRB 090510: $\xi \lesssim 0.1$ under the model assumptions adopted there [27,28].
- Fermi-LAT, GRB 130427A: $\xi \lesssim 0.5$ [29].
- H.E.S.S., PKS 2155-304 flare: $E_{QG,1} > 2.1 \times 10^{18}$ GeV, corresponding to $\xi \lesssim 6$ in the parametrization of Equation (24) [30].
- Cherenkov Telescope Array (CTA) expected reach: $\xi \sim 0.02$ [31].
- LHAASO, GRB 221009A: $E_{QG,1} > 10 E_P$ and $E_{QG,2} > 6 \times 10^{-8} E_P$ at 95% confidence level, providing the strongest published quadratic bound used here [32].

The geometric value $\xi_{\text{IVM}} = 0.577$ lies a factor of about six above the existing GRB 090510 bound and the LHAASO linear bound quoted above, although it remains compatible with the weaker H.E.S.S. blazar limit. Taken at face value, this naïve identification of the geometric ratio with a linear dispersion coefficient is therefore excluded by the stronger gamma-ray-burst constraints. There is also a theoretical reason to distrust a linear term: within effective field theory, linear CPT-odd modifications of dispersion relations arise from higher-dimensional operators and are subject to stringent low-energy constraints [33]. An order-unity linear coefficient is therefore also theoretically disfavoured.

It should be stressed that what these bounds exclude is a linear coefficient of order unity, such as the geometric value of Equation (60) would supply, and not a linear term

as such. The centrosymmetry theorem of Section 3.16 gives exactly zero for the linear coefficient, which satisfies each of the bounds above trivially. This section therefore records why an order-unity geometric identification of the linear coefficient is not tenable, and why the vanishing established in Section 3.16 is the only outcome consistent with the data.

3.16. The Quadratic Sector: Centrosymmetry, Angular Structure and Testability

The propagator of Section 3.13 settles both of the questions left open by the confrontation with data in Section 3.15. The vanishing of the observable linear coefficient is a theorem of the lattice construction rather than a symmetry motivation. The structure function of any real, symmetric nearest-neighbour quadratic form on a monatomic Bravais lattice is an even function of k: every odd moment of the neighbour distribution vanishes identically, the vanishing being exact rather than approximate, holding direction-by-direction rather than on the spherical average alone, and extending to every odd order in $\ell_0 k$. It follows that

$$\xi_{\mathrm{OBS}} = 0. \tag{76}$$

The generality of this statement should be emphasised: it is a property of the whole class of real nearest-neighbour Bravais lattices rather than a specific fingerprint of the FCC cell. What would generate a non-zero linear coefficient is a departure from that class: a complex or chiral hopping amplitude, a lattice with a non-trivial basis, or a CPT-odd operator of the Myers–Pospelov type, and none of these is present in the action of Equation (62). The prediction $\xi = 0$ is therefore a strong falsification channel, since any detection at any level would exclude the construction, but it is not by itself evidence for the particular lattice assumed here.

The dispersion relation is therefore governed at leading order by a quadratic term whose coefficient is computed rather than assumed. It is not a scalar but carries the angular dependence of the lowest octahedral harmonic. Propagation is therefore slowest along the body diagonals and fastest along the cube axes. Its spherical average is $1/20$, independent of direction and, for the diagonal equal-weight nearest-neighbour action of Equation (62), independent of the lattice coordination.

In particular, the value $1/3$ that would follow from identifying the coefficient with the square of the geometric projection factor of Equation (60) is larger than the computed isotropic part by a factor of about seven, and no identification of the lattice scale with the Planck scale that can be independently justified closes that gap. The geometric identification is therefore not tenable, and the coefficient must be taken from the propagator.

A single scalar coefficient is a weak signature, since a number of quantum-gravity phenomenologies can accommodate one. The octahedral angular pattern is fixed in form by symmetry, whereas its amplitude depends on the lattice and, more generally, on the choice of O_h-invariant couplings. For the diagonal equal-weight nearest-neighbour FCC action of Equation (62), the resulting normalised peak-to-peak modulation is:

$$\frac{\mathcal{A}_{\max} - \mathcal{A}_{\min}}{\langle \mathcal{A} \rangle} = \frac{5}{18} \simeq 27.8\% \tag{77}$$

The value $5/18$ is specific to twelve-fold FCC coordination. If the unitarity analysis instead selects eight-fold BCC coordination, the octahedral harmonic and angular ordering remain unchanged, but the normalised modulation becomes $20/27$, as derived in Section 3.13. For a cosmological source the propagation direction is fixed by the position of the source on the sky, so the signature is directional rather than temporal. A population of gamma-ray bursts or flaring blazars binned by sky position should show a modulation of the fitted time-delay coefficient of the amplitude given above, following the octahedral pattern of Figure 1. That pattern has six minima along the three coordinate axes and eight maxima along the

four body diagonals and therefore carries no dipole component, which distinguishes it from source-intrinsic systematics and from peculiar-motion contributions, both of which are predominantly dipolar. A sidereal modulation does arise, but in terrestrial interferometry rather than in astrophysical timing, since there the apparatus axis rotates with the Earth relative to the lattice frame; the two signatures should not be conflated.

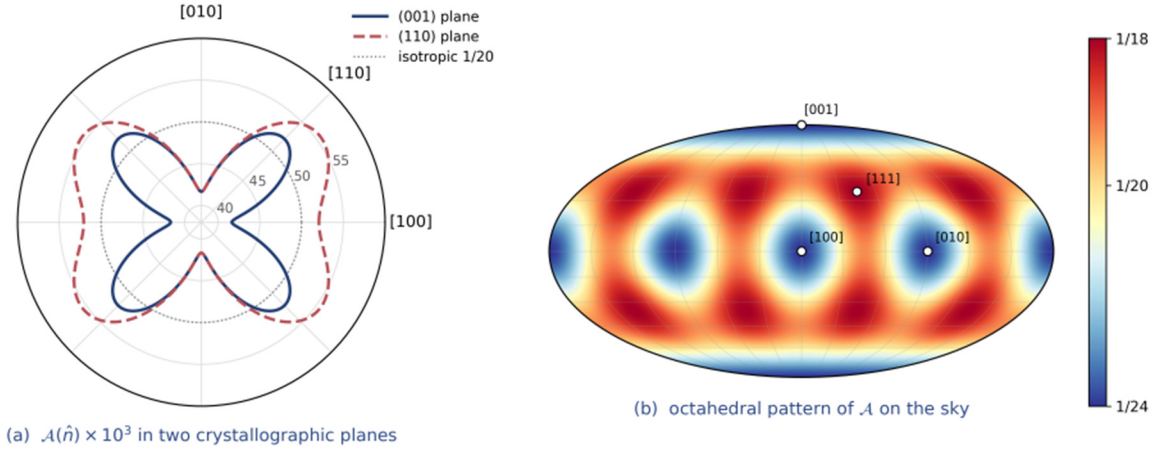


Figure 1. Angular structure of the quadratic dispersion coefficient in the twelve-fold FCC formulation. (**a**) Polar sections in the (001) and (110) planes, in units of 10^{-3}; the dotted circle is the isotropic value 1/20. (**b**) The same coefficient over the celestial sphere in Mollweide projection, with the high-symmetry lattice directions marked. Propagation is slowest along the body diagonals and fastest along the cube axes.

Two limitations of the prediction should be acknowledged here. The orientation of the lattice axes relative to any observable frame is not predicted by the framework, so the pattern is fixed in shape but free in orientation, and a fit against data therefore carries three additional free parameters. This weakens the test but does not remove it, since the shape has no dipole component and cannot be reproduced by an arbitrary reorientation. Furthermore, the calculation of Section 3.13 treats the gauge field as a set of components carried by the nodes; the polarisation structure is fixed by the symmetry argument of Section 3.17 rather than derived from a link-variable formulation, and a full lattice gauge treatment would be more satisfactory.

Quadratic Lorentz-violation bounds are substantially weaker than linear ones, and current observations remain compatible with a Planck-scale quadratic coefficient of the magnitude obtained in Section 3.13. The leading-order time-of-flight delay scales as

$$\Delta t \approx \frac{3}{2}\langle \mathcal{A} \rangle \frac{E_1^2 - E_2^2}{E_P^2} \int_0^z \frac{(1+z')^2}{H(z')}\, dz' \quad \underset{\text{flat, low-}z}{\rightarrow} \quad \frac{3}{2}\langle \mathcal{A} \rangle \frac{E_1^2 - E_2^2}{E_P^2} \frac{d}{c} \tag{78}$$

and is therefore extremely small. Using the isotropic coefficient of Equation (69) and the provisional normalisation $\ell_0 = \ell_P$, the delay is about 5×10^{-13}s for photons of 100 TeV over 1 Gpc, and about 5×10^{-21}s at 10 GeV. Expressed in the conventional parametrization,

$$\mathcal{A} = \left(\frac{E_P}{E_{QG,2}} \right)^2 \Longrightarrow E_{QG,2} \simeq 4.5 E_P, \text{ bound : } E_{QG,2} \gtrsim 6 \times 10^{-8} E_P. \tag{79}$$

The prediction corresponds to a quantum-gravity scale of roughly four and a half Planck energies. The strongest published quadratic time-of-flight bound used here, obtained by LHAASO from GRB 221009A, is $E_{QG,2} \gtrsim 6{\times}10^{-8} E_P$ at 95% confidence level, corresponding in the parametrization of Equation (79) to $|\mathcal{A}| \leq 3 \times 10^{14}$. The predicted coefficient $\langle \mathcal{A} \rangle = 1/20$ therefore lies approximately sixteen orders of magnitude below the current upper bound.

Even substantial future improvements in quadratic time-of-flight constraints would leave the predicted coefficient far below foreseeable time-of-flight sensitivity.

The quadratic coefficient is therefore not a falsifiable target for the Cherenkov Telescope Array or for any foreseeable time-of-flight measurement. What the data do provide, read in the opposite direction, is a bound on the lattice spacing:

$$|\mathcal{A}| = \frac{1}{20}\left(\frac{\ell_0}{\ell_P}\right)^2 \lesssim 3 \times 10^{14} \implies \ell_0 \lesssim 8 \times 10^7 \ell_P \sim 10^8 \ell_P. \tag{80}$$

This bound is genuine but extremely weak.

The same conclusion applies a fortiori to the octahedral anisotropy of Section 3.13, whose amplitude is 27.8 per cent of an already unobservable isotropic term. The sky pattern remains the geometric fingerprint of the framework, but it is not a measurement anyone can presently perform, and it should not be presented as one.

Where, then, is the framework falsifiable today? Two exact falsification channels follow from the discrete symmetries retained across the lattice constructions considered here. First, the linear coefficient vanishes identically: because the centrosymmetry theorem established above admits no small non-zero value, a detection of any linear time-of-flight term at any level would falsify the centrosymmetric vacuum. Second, and most sharply, helicity-odd birefringence is exactly absent, as established in Section 3.17: astrophysical polarimetry already constrains the helicity-odd coefficient at the 10^{-15} level and below [34], and CQST predicts that coefficient to vanish identically, so any positive detection would refute the inversion-symmetric lattice construction. In the current nearest-neighbour FCC formulation, the quadratic correction also has a definite subluminal sign in every direction. This sign is not, however, a formulation-independent prediction of CQST: as discussed in Section 4.6.2, introducing a four-dimensional custodial symmetry would alter the leading correction and require its sign to be re-derived. The two symmetry-based tests are therefore exact but asymmetric: they can refute the framework but cannot confirm it. The quadratic sector provides additional conditional tests of the present formulation, but its magnitude lies far below current observational sensitivity.

One point of consistency should be stated explicitly. The photon carries no rest metabolic rate; within the lattice constructions considered here, this is consistent with the exact vanishing of the linear propagation term, while photons can still probe the discrete geometry during propagation.

3.17. Absence of Vacuum Birefringence

Vacuum birefringence supplies the most stringent existing constraints on Lorentz-invariance violation, with gamma-ray-burst polarimetry providing particularly strong limits on helicity-dependent photon propagation [34,35]. A torsional geometry has, in principle, the potential to induce such effects. Since a birefringent framework would therefore be excluded immediately by existing observations; the question must be settled explicitly rather than passed over. The propagator of Section 3.13 settles it directly.

Birefringence requires the two photon helicity states to propagate with different phase velocities. In effective field theory, this effect can arise from CPT-odd operators of mass dimension five, characterised by contractions with the Levi–Civita symbol and hence by odd parity [36]. The vacuum polarisation obtained from the lattice action of Equation (62) has the form

$$\Pi_{\mu\nu}(k) = \left(\eta_{\mu\nu}k^2 - k_\mu k_\nu\right)F\left(k^2, K_4\right) + \Pi^{\text{aniso}}_{\mu\nu}(k),\ \epsilon_{\mu\nu\lambda\sigma} - \text{terms forbidden by } \mathcal{D} \in O_h. \tag{81}$$

The scalar part of the vacuum polarisation is governed by the function $F\left(k^2, K_4\right)$, while any parity-even polarisation-dependent contribution is contained in $\Pi^{\text{aniso}}_{\mu\nu}(k)$.

Two statements of different evidential status must be distinguished. The first is established. The octahedral group of the face-centred cubic lattice contains spatial inversion, so every parity-odd contraction vanishes identically. The CPT-odd operators of mass dimension five that generate helicity-dependent propagation are therefore excluded, and the predicted helicity-odd birefringent coefficient is exactly zero at this order. The same symmetry argument that removes the linear term in the dispersion relation excludes these operators as well, and the framework is consistent with gamma-ray-burst polarimetry bounds by construction rather than by adjustment.

The second conclusion does not follow from the present argument. Inversion symmetry does not forbid a parity-even structure capable of distinguishing two linear polarisations at a finite wave vector. This is not a hypothetical concern: cubic crystals of the very point group O_h exhibit measured birefringence induced by spatial dispersion at order k^2, an effect predicted by Lorentz and observed in CaF_2 and BaF_2 [37]. Its angular structure belongs to the same octahedral class as the harmonic K_4 computed in Section 3.13, while exhibiting a distinct directional dependence: it vanishes along $\langle 100 \rangle$ and $\langle 111 \rangle$ and peaks along $\langle 110 \rangle$. The tensor written in $\Pi^{\text{aniso}}_{\mu\nu}$ above is precisely the term that would carry such an effect, and it is not computed here: the calculation of Section 3.13 treats the gauge field as a set of components carried by the nodes, so the two transverse polarisations are degenerate by construction and the question cannot be settled within it.

Determining whether CQST predicts a parity-even birefringence, and with what octahedral pattern, requires the link-variable formulation listed in Section 4.6. On dimensional grounds any such effect would enter at order $(E/E_P)^2$, hence some 10^{-32} at TeV energies, far below any polarimetric sensitivity; it is therefore not in conflict with observation, but neither is its absence a result of the framework.

This conclusion depends on the treatment of torsion as a non-propagating auxiliary field adopted in Section 2.1. Were torsion instead a propagating field with a parity-odd coupling, the helicity-odd channel would reopen and the framework would be in immediate conflict with polarimetry. The status of the torsional sector and the absence of helicity-odd birefringence are therefore two aspects of a single question.

3.18. Residual Lorentz Violation and the Status of Boost Invariance

The microscopic construction rests on a discrete spatial lattice, and spatial isotropy alone does not guarantee full Lorentz invariance. Boost invariance in particular is not secured by any argument given in this paper, and the framework should not be read as claiming otherwise. This section states what is and is not established.

The lattice does generate preferred-frame effects after coarse-graining. The 3 + 1 split adopted in Section 2.6, in which the spatial lattice is primitive and the fourth dimension is generated by the tick, introduces a preferred foliation; this is the price of using the three-dimensional kissing number, and it is a real cost rather than a technicality. What the octahedral symmetry does provide is control over which Lorentz-violating operators can appear in the effective theory. Operators admissible under O_h invariance and under the discrete symmetries of the tick may be enumerated and compared with the standard Lorentz-violating operator classifications employed in phenomenological analyses [38].

The enumeration is short and worth giving explicitly (Table 4). Operators are classified by mass dimension and by their behaviour under the octahedral group and under the discrete symmetries of the tick.

Table 4. Symmetry status of the principal Lorentz-violating operators.

Operator/Effect	Status in CQST	Reason
Dimension-4 operator $c^{ij}F_{i\mu}F_j^{\mu}$	Isotropic by symmetry	The only O_h-invariant rank-two tensor is δ^{ij}, so any such term is isotropic and is reabsorbed into the normalisation of c; no anisotropy is possible at this order.
CPT-odd dimension-5 Myers-Pospelov operator	Excluded by inversion symmetry	Odd under spatial inversion; O_h contains inversion
Linear dispersion term $\xi(E/E_P)$	Excluded by centrosymmetry	Vanishes by centrosymmetry of the neighbour set
Helicity-odd birefringence	Excluded by inversion symmetry	Requires parity-odd operator
Quadratic isotropic term $\mathcal{A}(E/E_P)^2$	Derived in the present formulation	Generated by the propagator of Section 3.13
Octahedral anisotropy $K_4(\hat{\boldsymbol{n}})$	Derived in the present formulation	Lowest non-trivial harmonic compatible with O_h
Parity-even birefringence from spatial dispersion	Not computed	Would require the tensor structure $\Pi^{\text{aniso}}_{\mu\nu}$ and a link-variable formulation
Exact Lorentz invariance	Not established	Preferred foliation remains
Dimension-6 Lorentz-violating operators	Allowed by symmetry	Planck-scale suppression
Radiative stability of the Planck suppression	Not addressed	No custodial symmetry is available in the present formulation; candidate forms are indicated in Section 4.6.2

The absence of anisotropy at dimension four is a group-theoretic consequence of the octahedral symmetry. The space of spherical harmonics with $\ell = 2$ contains no A_{1g} representation of O_h, whereas the $\ell = 4$ sector contains exactly one (Appendix A, Appendix A.4). Anisotropy therefore appears only at quartic order, making the result of Section 3.13 a consequence of symmetry rather than merely a feature of the explicit calculation.

CPT-odd dimension-five operators are excluded by the octahedral symmetry, as shown in Section 3.17. Their suppression at low energies is by Planck-scale power counting, which is a statement about magnitude and not a mechanism guaranteeing exact boost invariance.

This is not a formality. It has been shown that Lorentz-violating operators suppressed by the Planck scale generically induce dimension-four Lorentz violation with coefficients of order unity through radiative corrections, unless a custodial symmetry protects against it [39]. CQST possesses no such symmetry in its present formulation, and the naturalness problem this raises is not solved here. It is recorded among the open problems of Section 4.6 rather than passed over, since the phenomenological content of Section 3.16 rests on dimension-six Lorentz-violating operators of the type to which the argument applies. The form such a symmetry might take, and its consequences for that phenomenological content, are discussed in Section 4.6.2.

A remark on the alternative route is in order. Poisson sprinkling, as employed in causal-set theory, restores Lorentz invariance in a well-defined statistical sense, a property that can be understood as a consequence of the compatibility between random causal-set discreteness and Lorentz symmetry [40]. It has not been adopted, because it would destroy the very geometric regularity from which the dispersion signature of Section 3.16 is derived: the framework would become Lorentz-invariant at the cost of losing its principal

phenomenological signature. The same requirement applies to the irregular configurations needed for the gravitational sector: the dominant phase must retain the orientational order of Ω_0 over cosmological baselines if the octahedral signature is to survive (Section 4.6.1). The position taken here is the opposite of the sprinkling approach. The residual anisotropy is retained, its unresolved status at the level of exact boost invariance is stated openly, and it is presented as an observational signature rather than concealed. CQST accordingly predicts a specific, calculable and testable pattern of Lorentz violation; it does not claim to recover exact Lorentz invariance, and the recovery of boost invariance in the continuum limit remains an open problem listed in Section 4.6.

3.19. The Extended Vacuum Fixed Point

The vacuum fixed point Ω_0 is enlarged to include the transport sector, the scalar order parameter, the vacuum causal density and the infrared electromagnetic coupling. The corresponding extended fixed-point map is denoted by $\mathcal{F}_{\text{ext}}$. The extended vacuum fixed point is characterised by the following configuration and fixed-point conditions:

$$\Omega_0 = \left(U*_{ij}, \rho*, C_0*, \alpha*\right), \tag{82}$$

$$\text{with } \mathcal{F}_{\text{ext}}[\Omega_0] = \Omega_0, \tag{83}$$

and the characteristic infrared relation (a one-parameter relation in b; see Section 3.5)

$$\alpha^* = \frac{1}{\left[\left(\frac{b}{2\pi}\right) ln\left(\frac{m_P c^2}{\hbar H_\infty}\right)\right]}, \tag{84}$$

as in Equation (33) and

$$b \approx 6.13 \rightarrow \alpha* \approx 1/137. \tag{85}$$

The full set of fixed-point outputs is then proposed to include the emergent metric $g_{\mu\nu}$, the vacuum causal density C_0, the infrared coupling α, a positive cosmological constant Λ, Lorentzian signature, and the dispersion structure discussed in Section 3.13. At the present stage, however, this list should be understood as a statement of scope rather than of derivation: C_0 enters through its definition in terms of Λ(Section 2.7, Section 3.6), whereas α appears through the one-parameter calibration discussed in Section 3.5. Ω_0 in (82) denotes the complete vacuum fixed-point configuration, whereas ρ* denotes the equilibrium value of the scalar order parameter within that configuration. The geometric vector equilibrium discussed in Section 3.1 is interpreted as the geometric realisation of Ω_0 rather than as an independent definition of the fixed point itself.

4. Discussion

4.1. The Derivative Chain

The conceptual synthesis is the chain: $\hbar$ as the cost of a causal tick $\rightarrow E = h\nu$ for oscillatory modes; the Compton frequency as a tick rate (Equation (44)); rest-energy bookkeeping $\rightarrow E = mc^2$; the minimum cost of observation $\rightarrow$ the uncertainty relations; the de Broglie relation (Equation (52)); and the finite propagation time of one tick $\rightarrow c$ as the maximal microscopic causal speed. CQST thus replaces several separate postulates with one microscopic principle—the finite, discrete, energetically non-zero cost of an elementary causal relation—while leaving the mathematical rigour of this replacement as open work.

The epistemic status of this chain should be read alongside Table 3. Its links are of two kinds, and they are not interchangeable. The identification of $\hbar$ with the cost of a causal tick, and the relations that follow from it—the energy–frequency relation, the mass–frequency correspondence, rest-energy bookkeeping, the uncertainty relations and the recovery of

Schrödinger dynamics—are reinterpretations of relations established independently of CQST. They demonstrate the internal consistency of the causal-tick reading and constitute no evidence for it. Only the final link, from the discrete lattice to the photon propagator of Section 3.13, is a derivation in the strict sense, and only it produces statements that could turn out false.

4.2. Unified Logical Map

The central logical dependencies are listed below. Solid arrows denote implications established within the present work, whereas broken arrows denote assumptions or conjectured links whose derivation has not yet been obtained:

- Discrete tick operator $\hat{\Theta}_{\mathrm{QSR}}$ and Hamiltonian $\hat{H}_{\mathrm{QSR}}$ ⇢ effective continuum description through the coarse-graining scheme of Section 2.10 (assumed, not derived; see Section 2.10).
- Causal graph $G = (V, L)$ ⇢ emergent metric $g_{\mu\nu}$ through the correlator ansatz of Equation (4) (assumed, not derived; Section 2.1).
- IVM geometry, maximal valence 12 ⇢ conjectural geometric reading of the one-quarter entropy factor.
- Causal acyclicity ⇢ Lorentzian signature and holographic scaling.
- Cosmological constant $\Lambda \to$ vacuum causal density C_0 ⇢ calibrated infrared value of α (no epoch dependence; Section 4.5.2).
- Finite energetic cost of a causal tick ⇢ quantum relations reinterpreted in terms of $\hbar$ (Sections 2.8 and 3.8–3.11).
- IVM lattice structure → computed quadratic dispersion coefficient $A(\hat{n})$ via the photon propagator (Section 3.13).
- Vacuum fixed point Ω_0 ⇢ proposed common source of geometry, couplings and phenomenology (Section 3.19).

4.3. Effective Local Response and Emergent Gravity: A Speculative Sector

A further speculative possibility does not modify the fundamental quantum of action $\hbar$, which remains the fixed cost of a causal tick as defined in Equation (23). Instead, spatial variations in the local causal density $C(x,t)$ may generate effective modifications of the coarse-grained response of the lattice. Such effects, if present, would arise from gradients of the causal-density field rather than from variations in $\hbar$ itself.

The cosmological quantity C_0 introduced in Section 2.7 represents the homogeneous background causal density, whereas $C(x,t)$ denotes local departures from that background. In regions where $C(x,t)$ varies spatially, the statistical properties of causal transport may differ slightly from those of the vacuum state Ω_0. Such effects could, in principle, produce effective forces or trajectories governed by gradients of the causal-density field,

$$\nabla C(x,t) \to \text{effective gravitational response.} \tag{86}$$

In this interpretation, free-fall behaviour would not arise from variations in the fundamental quantum of action, but from the tendency of coarse-grained causal flows to evolve along paths determined by the local structure of the causal-density field (see Section 2.7 for the distinction between the homogeneous background causal density C_0 and the local field $C(x,t)$). In the present section, only gradients of $C(x,t)$, not variations in the fundamental quantum of action $\hbar$, are assumed to contribute to the speculative effective gravitational response.

We emphasise that this idea remains highly speculative. At present, no complete dynamical derivation connects gradients of $C(x,t)$ to the Einstein gravitational field, and the proposal should therefore be regarded as a heuristic possibility rather than an established

prediction of CQST. The central framework developed in Sections 3.5–3.16 does not depend on this hypothesis.

4.4. Predictions from a Causal-Density-Dependent Effective Quantum Response

If local causal-density variations modify the effective coarse-grained quantum response of the lattice, the quantum sector becomes weakly environment-dependent. Three classes of consequence follow, all small and at present unobserved: tiny effective deviations from standard quantum dynamics in gravitational fields, negligible near Earth but potentially amplified near compact objects; additional spectral shifts in strong gravity beyond the standard gravitational redshift; and small effective-mass shifts in regions of high causal density, visible only in very high-precision spectroscopy. These are offered as constraints-to-be-saturated rather than as confirmed signatures.

4.5. Additional Observational Consequences

Other potential signatures include dispersion or polarisation effects in the gravitational-wave sector, on which Section 4.5.1 comments. Two signatures that might be expected of a framework of this kind warrant separate discussion. A helicity-dependent rotation of photon polarisation would require a parity-odd birefringent term. Section 3.17 shows that such helicity-odd terms are excluded by spatial inversion. Parity-even birefringence induced by spatial dispersion is not excluded by the point group and remains to be computed.

4.5.1. The Gravitational-Wave Sector

The gravitational-wave implications discussed here are necessarily conditional because the gravitational sector and its propagating tensor modes have not yet been derived within CQST. Assuming that spin-2 tensor modes propagate on the same O_h-invariant substrate as the photon sector, the symmetry argument of Section 3.17 applies: helicity-odd birefringence of a massless field requires a parity-odd structure, and the inversion element of the octahedral point group therefore excludes helicity-odd gravitational birefringence. Parity-even birefringence induced by spatial dispersion is not excluded by symmetry and remains to be computed. Under the same assumption, chromatic dispersion of gravitational waves is permitted: it is even under parity and may display the same quadratic dependence as in the photon sector. Such an effect would, however, be entirely unobservable at detectable gravitational-wave frequencies. The gravitational-wave sector therefore provides only a conditional falsification channel: a detection of helicity-odd gravitational birefringence would refute the inversion-symmetric propagation hypothesis without appeal, there being no parameter available to adjust.

4.5.2. The Epoch Dependence of α: A Quantitative Conflict

If the causal density is allowed to evolve with H while α remains a function of that density, the framework predicts a redshift evolution of the fine-structure constant. As shown below, this prediction is in quantitative conflict with observation. That prediction has been tested, repeatedly and with high precision, and it fails by a wide margin. This section derives the conflict rather than asserting it, and shows that the definition adopted in Section 2.7 avoids it structurally rather than by stipulation.

Suppose, contrary to Section 2.7, that the causal density were defined through the instantaneous expansion rate, so that the relation of Section 3.5 held at each epoch with $H \to H(z)$. Differentiating with respect to time and using the deceleration parameter at the present epoch gives

$$\frac{\mathrm{d}\alpha^{-1}}{\mathrm{d}t} = \frac{b}{2\pi}(1+q_0)H_0 \simeq 3.0\times10^{-11}\mathrm{yr}^{-1} \Rightarrow \ |\frac{\dot{\alpha}}{\alpha}| \simeq 2\times10^{-13}yr^{-1} \tag{87}$$

with $b/2\pi = 0.976$ and $q_0 \simeq -0.55$. Comparisons of optical atomic clocks, particularly measurements of the $^{171}Yb^+$ E3/E2 frequency ratio, constrain the fractional drift of the fine-structure constant at the level of 10^{-18} yr^{-1} [41]. The drift predicted under the epoch-dependent definition of the causal density therefore exceeds the laboratory bound by five to six orders of magnitude. The corresponding cosmological variation follows by integration,

$$\frac{\Delta\alpha}{\alpha}(z) = \frac{b}{2\pi\alpha^{-1}} ln \frac{H(z)}{H_0} \simeq 5 \times 10^{-4} (z \simeq 0.14), \;\; 8 \times 10^{-3} \; (z \simeq 2) \tag{88}$$

for a standard ΛCDM (Lambda Cold Dark Matter) expansion. The Oklo natural reactor constrains the fractional variation below 10^{-7} at the lower redshift [42], while high-precision quasar absorption spectroscopy constrains variations in α at the parts-per-million level at higher redshift [43]; the predicted variation therefore remains several orders of magnitude above both classes of constraint. This is decisive against the epoch-dependent reading, and it is worth being clear about which of the two ingredients fails. The coefficient b is not the culprit. In a genuine field theory b is not a constant—it is the coefficient of a beta function and changes at every mass threshold—but it runs with energy scale, and the thresholds sit at particle masses, which do not move with cosmic epoch. Making b a free function of redshift would in any case destroy the relation: with $b(z)$ unconstrained, any history of α can be reproduced, and a one-parameter fit becomes a fit to a function. Moreover, the cancellation would have to be fine-tuned to two parts in ten thousand to respect the Oklo bound, which is not a small correction but an exact conspiracy imposed by hand.

What fails is the other ingredient, the identification of the vacuum causal density with the instantaneous expansion rate. $H(z)$ is not a constant of the spacetime: it contains the matter density, which dilutes, and decays monotonically towards the asymptotic value fixed by Λ alone. There is no evident reason why a property of the vacuum should track the amount of matter present at a given epoch. Defining the causal density through Λ instead, as Section 2.7 does, gives

$$\frac{dC_0}{dt} = 0 \implies \frac{\dot{\alpha}}{\alpha} = 0 \;\text{ identically} \tag{89}$$

and the conflict does not arise. The cost is negligible and the gain is structural. Numerically, nothing of consequence changes: the asymptotic rate is $\Omega_\Lambda^{(1/2)} \simeq 0.83$ times the present one, the logarithm shifts by 0.19 in 140, and the calibrated coefficient moves from 6.14 to 6.13—immaterial, since b is fitted in either case. What is gained is that the relation now holds between two constants of nature rather than between a constant and an evolving quantity, which is what made it ill-posed before.

Three qualifications are owed to the reader. First, this is a reason and not a derivation: nothing in the framework compels the definition through Λ, and it is adopted because it is the only reading under which the relation is well posed. Second, it does not improve the evidential standing of the relation, which remains the one-parameter calibration of Section 3.5; removing an inconsistency is not the same as gaining support. Third, the argument assumes Λ to be a true constant. Were the dark energy dynamical, the causal density would inherit that evolution and a weakened form of the conflict would return, bounded by the observational limits on the dark-energy equation of state. Under that reading the constancy of α places an additional constraint on the allowed evolution of dark energy. This is a conditional consequence of the dynamical-dark-energy extension, rather than an independent test of the constant-Λ formulation adopted here.

A search for redshift evolution of α is accordingly not among the experimental priorities of Section 4.7. It is not a test of CQST under the definition adopted here, which predicts no evolution at all; the observed constancy of α is instead one of the considerations that fixes the definition.

4.6. Limitations and Theoretical Priorities

The CQST framework brings together geometric, holographic, quantum-foundational, and phenomenological ideas within a common causal-spacetime picture. The next stage of development is to deepen the mathematical structure of the theory, clarify its foundational mechanisms, and strengthen its predictive capacity. The theoretical priorities listed below outline the principal directions for advancing the framework and establishing more explicit connections between its underlying causal dynamics and observable physical phenomena:

- Derive the continuum limit. The effective Lagrangian of Section 2.1 is at present an ansatz constrained by the symmetries of the substrate, not a consequence of the discrete dynamics of Section 2.3. A controlled derivation would proceed either through a path integral over causal configurations or through an explicit renormalization procedure on the graph, and would have to reach the gravitational sector: no curvature term is constructed here and the passage from the discrete dynamics to the Einstein equations is the principal missing element of the framework. The non-perturbative difficulties this entails, and their relation to Causal Dynamical Triangulations and other discrete approaches, are discussed in Section 4.6.1.
- Identify the unitarity condition on the coin operator, restricted to the T_{1u} channel, that obstructs twelve-fold coordination. Appendix A shows that the obstruction cannot originate in the symmetry content of the bond space, because the vector representation T_{1u} occurs uniquely in both the eight-channel and the twelve-channel spaces.
- Make the one-quarter entropy factor precise by counting the independent boundary degrees of freedom of the discrete lattice.
- Reformulate the photon sector in terms of link variables $U_{ij} = \exp\left[\frac{ie}{\hbar}\int_{x_j}^{x_i} A_\mu dx^\mu\right]$ and a plaquette action, thereby making gauge invariance manifest at finite lattice spacing and deriving the polarisation structure, including the parity-even birefringence left unresolved in Section 3.17.
- Determine whether the QSR substrate possesses a custodial symmetry capable of protecting the Planck suppression of the Lorentz-violating operators against radiative promotion to dimension four. A candidate form of such a symmetry, and its consequences for the phenomenology of Section 3.16, are indicated in Section 4.6.2.
- Clarify whether the conformal response parameter β can be derived from the microscopic dynamics, despite carrying no predictive content in the present formulation. Although the conformal dressing $g_{\mu\nu} = e^{-\beta C_0}\tilde{g}_{\mu\nu}$ can be written formally, it is not used in the present formulation of CQST. Since $C_0 \sim 10^{-122}$, any finite value of β produces an exponentially small correction and therefore no phenomenologically relevant effect. The parameter β should consequently be regarded, at this stage, as a purely speculative coarse-grained response coefficient of the QSR lattice. Future lattice-based analyses could investigate whether a non-trivial effective value emerges from the microscopic dynamics. Until such a derivation is available, β carries no predictive content and plays no role in the observational results discussed in this paper.
- Compute the coefficient b from an explicit treatment of the Standard Model sector, turning the α relation from a fit into a prediction.
- Establish the existence and stability of Ω_0 in an appropriate functional space (e.g., a Banach–Schauder-type argument).
- Determine whether a causal-density-dependent effective quantum response can be made consistent with unitarity, effective field theory, covariance, and precision tests of quantum mechanics.
- Reconcile the cosmological holography with a positive cosmological constant, since de Sitter holography lacks the controlled status of AdS/CFT.

4.6.1. The Continuum Limit in the Context of Non-Perturbative Quantum Gravity

Extracting a four-dimensional continuum geometry from a discrete state sum is a hard non-perturbative problem, and the first of the priorities listed above should be read against the experience of other discrete approaches to quantum gravity. In Euclidean dynamical triangulations the sum is dominated by degenerate geometries, crumpled configurations of very large Hausdorff dimension and branched polymers of Hausdorff dimension two, separated by a first-order transition, so that no continuum limit with an extended four-dimensional geometry is obtained [6,44,45]. Causal Dynamical Triangulations changed the picture by imposing a proper-time foliation and forbidding changes in spatial topology: the state sum then possesses a phase in which an extended four-dimensional geometry emerges [3,46], whose volume profile reproduces that of Euclidean de Sitter space [47], and second-order transition lines have been located at which a continuum limit might be defined [48]. Even there, however, the existence of that limit has not been established; the naive Regge curvature does not survive it and has had to be replaced by quantum Ricci curvature, and what has been recovered are global and averaged geometric observables rather than the Einstein equations themselves [49,50]. In causal-set theory the same obstacle appears in entropic form: non-manifold-like orders overwhelmingly dominate the sum, and their suppression by the causal-set action has been demonstrated only for particular classes [5,51].

The QSR substrate shares with CDT the features that proved decisive there, namely a causal order, a foliation generated by the tick and a cell complex of simplicial type on which curvature can be carried by Regge deficit angles [52], and it therefore inherits the same questions. The regular IVM complex is flat, since two tetrahedra and two octahedra close exactly around every edge, $2\arccos(1/3) + 2\arccos(-1/3) = 2\pi$; curvature must reside in departures from this closure, and the gravitational sector can arise only from a sum over irregular configurations, for which the regular lattice can serve at most as the vacuum configuration Ω_0. The octahedral dispersion pattern of Section 3.16 therefore presupposes that the dominant phase retains the orientational order of Ω_0 over cosmological baselines; whether this occurs in a state sum over irregular configurations is part of the unresolved continuum-limit problem.

Deriving the Einstein equations accordingly requires showing, in order, that a state sum over causal QSR configurations possesses a phase of extended four-dimensional geometry rather than a degenerate one, that this phase is bounded by a second-order transition at which a continuum limit can be taken, and that the effective dynamics of the emergent geometry is that of the Einstein–Hilbert action with a positive cosmological constant. The experience of CDT indicates, moreover, that curvature should be probed through coarse-grained observables rather than through raw deficit angles, and its de Sitter volume profile offers a concrete intermediate benchmark, particularly apt here since C_0 is defined through Λ (Section 2.7). The foliation is a further point of contact: in CDT it has been related to Hořava–Lifshitz gravity [53,54], so that the continuum limit of a foliated state sum need not be general relativity, which ties this priority directly to the status of boost invariance discussed in Section 3.18.

4.6.2. The Form of a Custodial Symmetry

The obstruction to radiative stability is group-theoretic, and this allows the form of a custodial symmetry to be indicated with some precision. The octahedral group leaves the tick direction u invariant, so that an operator such as $(u\cdot\partial)^2$ is O_h-invariant: the point group makes the dimension-four term isotropic, as Table 4 records, but it cannot make it universal, and a species-dependent limiting speed is precisely what radiative corrections generate [39]. A custodial symmetry must therefore act on the four-dimensional causal complex and

move the tick direction. The natural candidate is the D_4 lattice already encountered in Section 2.6 [55]. When the tick is chosen along a coordinate axis of D_4, the orthogonal spatial slice contains the twelve cuboctahedral vectors of the IVM. The remaining twelve nearest-neighbour vectors separate into six future-directed and six past-directed displacements of half the cubic cell side along the cube axes, which become null in Lorentzian signature. Here $d = a/\sqrt{2}$ is the nearest-neighbour separation within the FCC spatial slice, whereas each null inter-slice link has spatial projection $a/2 = d/\sqrt{2}$ and temporal separation $c\tau_4 = a/2$. Thus, the scale $\ell_0 = d$ used in Equations (70) and (80) denotes the spatial nearest-neighbour separation in the original FCC construction and is distinct from the spatial projection $a/2 = d/\sqrt{2}$ of each inter-slice link in the D_4 extension. The coordination twelve of the IVM and the four-dimensional kissing number twenty-four of Section 2.6 are then two aspects of one structure, and O_h is precisely the stabiliser of the chosen tick direction in the Weyl group $W(F_4)$, of order 1152.

The hypercubic subgroup $W(B_4)$, which permutes the tick with the three cube axes, is the minimal custodial group: its only invariant symmetric tensor of rank two is the metric, so that a common limiting speed is enforced for every species and cannot be undone by radiative corrections. This is how Euclidean lattice gauge theory obtains a single speed of light without tuning [26], whereas on anisotropic lattices, where no symmetry relates the temporal to the spatial directions, the gauge and fermion anisotropies must be tuned separately [56]. The full group $W(F_4)$ protects further: its invariant tensors of rank two and four are exactly those built from the metric, excluding Lorentz-violating operators of dimensions four and six in the non-birefringent propagation sector. Dimension-six birefringent structures are not covered by this argument and remain unresolved, as noted in Section 3.17. Protection at every order would need an infinite symmetry group. A lattice admits one only in Lorentzian signature, as an arithmetic subgroup of the Lorentz group, which by the Borel density theorem has exactly the invariant tensors of the full group [57]; however, its orbits of timelike and null vectors are infinite, so that an invariant action cannot couple a node to finitely many neighbours. This is the trade-off between discreteness, locality and Lorentz invariance that leads causal-set theory to random sprinkling [40].

The imposition of such a symmetry would modify the phenomenology of Section 3.16, rather than merely ensure its radiative stability.

Under $W(F_4)$ the non-birefringent dimension-six operators on which Section 3.16 rests are excluded together with the dimension-four ones: the quadratic coefficient vanishes identically and the leading non-birefringent correction is quartic in E/E_P, with a sign not fixed by symmetry; for the nearest-neighbour D_4 analogue of Equation (62), direct expansion gives a group velocity exceeding its low-energy limiting value c away from the cube axes, although the computation is not shown here. This does not imply propagation beyond one elementary causal link per update, but shows that the microscopic causal bound alone does not determine the sign of the effective dispersive correction.

Under $W(B_4)$ a quadratic term survives only through couplings beyond nearest neighbours, with the nearest-neighbour D_4 action being automatically $W(F_4)$-invariant; its amplitude and sign are then set by those couplings, while its angular dependence is fixed by the group as $1 + (5/8)K_4$, with an ordering opposite to that of Equation (67). The exclusion of the linear term and of helicity-odd birefringence would survive, since both groups contain spatial inversion, but the sign of the leading correction would have to be re-derived. The structural price is that the tick must become a lattice direction on the same footing as the spatial links. This requires a Euclidean continuation of the tick dynamics with a positive transfer matrix, and it replaces the primitive $3 + 1$ split of Section 2.6 by a foliation chosen within a four-dimensionally symmetric complex, which would bear directly on the

preferred-frame problem of Section 3.18. Whether the unitary tick of Section 2.3 admits such a continuation is itself an open question.

4.7. Experimental Priorities

Having outlined the theoretical limitations and priorities of CQST, we now identify the observations capable of refuting the exact symmetry predictions of the framework or of constraining its conditional sectors.

- Continue to tighten bounds on a linear time-of-flight coefficient with Fermi-LAT, CTA and related observatories. The centrosymmetry theorem of Section 3.16 predicts exactly zero, so any positive detection would refute the centrosymmetric lattice construction; the quadratic coefficient itself is not a realistic target, for the reasons given there.
- Improve X-ray and gamma-ray polarimetry of GRBs and blazars. CQST predicts exactly zero helicity-odd birefringence, so any positive detection would refute the inversion-symmetric lattice construction (Section 3.17).
- Test whether the dark energy is a true cosmological constant. Under the definition of Section 2.7 the framework predicts no evolution of α whatever; a dynamical-dark-energy component would reintroduce a weak epoch dependence, so improving limits on the equation of state constrains CQST indirectly through the observed constancy of α.
- Search for helicity-odd gravitational birefringence. Assuming that spin-2 tensor modes propagate on the same O_h-invariant substrate as the photon sector, inversion symmetry predicts its exact absence; any positive detection would refute that propagation hypothesis (Section 4.5.1). The chromatic gravitational-wave dispersion conditionally permitted under the same assumption is far below observational sensitivity.
- Use precision spectroscopy near compact objects to bound causal-density-dependent effective quantum responses, including possible shifts in local Compton frequencies.

5. Conclusions

CQST joins two strands into one framework. The first is geometric: spacetime emerges from a discrete causal graph whose maximally symmetric realisation is identified with Fuller's Isotropic Vector Matrix. The second concerns quantum foundations: the reduced Planck constant is reinterpreted as the energetic cost of a causal tick of that same lattice. Together they attempt to reduce a broad range of physical structures to a single regenerating causal substrate, with the vacuum fixed point Ω_0 serving as the common source of the large-scale (de Sitter) metric, the causal density C_0, the cosmological constant, the infrared value of α, and a dispersion invariant.

The quantitative core of CQST must be stated narrowly. The relation between α and the causal density is a one-parameter fit and reproduces the observed coupling by construction; the cosmological relations are reparametrisations of quantities already defined in terms of Planck and de Sitter scales; and the holographic scaling and its one-quarter coefficient are conjectures. None of these is a prediction, and Table 3 records each as such.

The empirical content of the framework lies in a single place. The discrete electromagnetic action on the FCC lattice yields a photon propagator whose linear term vanishes by symmetry and whose leading quadratic correction, in the current nearest-neighbour FCC formulation, is subluminal in every direction and anisotropic, carrying the lowest octahedral harmonic. The octahedral form of that angular correction is fixed by symmetry, but its relative amplitude depends on the lattice action: the coefficient 5/12 follows specifically from the diagonal equal-weight nearest-neighbour FCC action of Equation (62) and may be modified by more general O_h-invariant couplings. Its overall magnitude is also not fixed by geometry alone, since it depends on the normalisation relating the lattice scale to the Planck

length. It should also be emphasised that the identification of $\hbar$ with the causal tick fixes its physical interpretation rather than its numerical value; the standard quantum relations are recovered as bookkeeping consequences of the causal dynamics rather than derived from a more fundamental constant.

CQST is at present better described as a speculative but testable research framework than as a developed fundamental theory. Important ingredients are missing: a controlled continuum limit, the microscopic determination of b, the recovery of exact boost invariance, a custodial mechanism ensuring the radiative stability of the Planck suppression (Section 4.6.2), and a resolution of the tension between the assumed twelve-fold coordination and the unitarity conditions on the tick operator.

What the framework does offer is a calculated photon sector with two exact, symmetry-based statements: the linear coefficient vanishes by centrosymmetry, and the helicity-odd coefficient vanishes because of the inversion element of the octahedral group. In the current nearest-neighbour FCC formulation, the leading quadratic correction is subluminal in every direction and, for the diagonal equal-weight action of Equation (62), has an isotropic coefficient of $1/20$ and an octahedral dependence characterised by the harmonic coefficient $5/12$, corresponding to a normalised peak-to-peak modulation of $5/18$ for twelve-fold coordination. These numerical coefficients are specific to the stated action, while the FCC result remains conditional on the unitarity question discussed in Section 2.6.

The first two provide exact falsification channels, whereas the quadratic result is both formulation-dependent and beyond current observational reach: its predicted magnitude lies approximately sixteen orders of magnitude below existing bounds. These are asymmetric tests, capable of refuting the framework but not of confirming it. That is a narrow evidential base, and we prefer to state it than to overstate it. The thesis itself remains coherent: that continuum geometry and quantum law are both emergent consequences of a discrete causal metabolism, and that the geometry of closest packing leaves a specific, calculable signature in the propagation of light.

Funding: This research received no external funding.

Data Availability Statement: No external datasets were used in this study. The numerical calculations reported in Section 3.13, the results displayed in Figure 1 and the computational checks described in Appendix A are reproducible from the analytical expressions and algorithms specified in the manuscript.

Acknowledgments: During the preparation of this manuscript, AI-assisted language and editorial tools (i.e., Claude Fable 5.1 and Microsoft 365 Copilot [GPT-5]) were used to support linguistic refinement, consistency checks and formatting review. The author reviewed, edited and validated all AI-assisted outputs and takes full responsibility for the scientific content of the publication.

Conflicts of Interest: The author was employed by RSE—Ricerca sul Sistema Energetico SpA. The author declares no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AdS	Anti de Sitter
BCC	Body-centred cubic
CDT	Causal dynamical triangulations
CFT	Conformal field theory
CQST	Causal quantum spacetime theory
CTA	Cherenkov telescope array
ECSK	Einstein–Cartan-Sciama-Kibble
FCC	Face-centred cubic

GRB	Gamma-ray burst
IVM	Isotropic Vector Matrix
LIV	Lorentz-Invariance Violation
QSR	Quantum spacetime regenerating
RG	Renormalization group
SC	Simple cubic

Appendix A. The Twelve-Channel Bond Space of the IVM Cell

This appendix carries out the finite computation posed in Section 2.6 and reports its outcome. That section asked whether the twelve transport channels of the local IVM cell constitute a single orbit of the automorphism group of the cell, and how the twelve-dimensional bond space V_{12} of Equation (16) decomposes into irreducible representations of it. Both questions are answered here. The first answer is affirmative, the second is explicit, and the consequence for the uniqueness question is negative in the precise sense that Section 2.6 anticipated: the decomposition identifies the face-centred cubic cell but does not select it. Coordination twelve therefore remains a defining assumption of the framework rather than a derived consequence, and is classified accordingly in Table 3.

Three further results follow from the same computation and are of independent use to the body of the paper. The first is a justification of the degree matrix asserted in Equation (16). The second is to establish the group-theoretic origin of the isotropic leading term and the uniqueness of the leading anisotropy discussed in Section 3.18. The third is an explicit statement of how much freedom the point group leaves in the choice of the lattice action of Equation (62). Every statement below has been verified by exhaustive computation; the checks are listed in Appendix A.8.

Appendix A.1. The Local Cell and Its Channel Graph

Let a QSR node sit at the origin of the face-centred cubic lattice and let a be the side of the conventional cubic cell. The twelve nearest-neighbour vectors are those of Equation (61),

$$\{\boldsymbol{e}_i\}_{i=1}^{12} = \left\{\tfrac{a}{2}(\pm 1, \pm 1, 0), \tfrac{a}{2}(\pm 1, 0, \pm 1), \tfrac{a}{2}(0, \pm 1, \pm 1)\right\}, \tag{A1}$$

with nearest-neighbour separation $d = a/\sqrt{2}$. These are the twelve vertices of a cuboctahedron. Declare two channels adjacent when the corresponding vertices are joined by an edge of that solid, that is, when they subtend an angle of sixty degrees:

$$\boldsymbol{e}_i \sim \boldsymbol{e}_j \iff \boldsymbol{e}_i \cdot \boldsymbol{e}_j = \tfrac{a^2}{4} \iff \left\|\boldsymbol{e}_i - \boldsymbol{e}_j\right\| = d = \tfrac{a}{\sqrt{2}}. \tag{A2}$$

The resulting channel graph Γ_{12} has twelve vertices and twenty-four edges and is four-regular: each nearest-neighbour vector stands at sixty degrees from exactly four others. The degree matrix of Equation (16) is therefore fixed, and the combinatorial Laplacian of the cell is

$$L = D - A, \quad D = 4I_{12}, \quad A_{ij} = \begin{cases} 1, & \mathbf{e}_i \sim \mathbf{e}_j, \\ 0, & \text{otherwise.} \end{cases} \tag{A3}$$

This is the justification of the value $D = 4I_{12}$ written in Equation (16). It is a property of the cuboctahedron rather than a choice, and the sixty-degree coordination that fixes it is the same one invoked in Sections 2.4 and 2.5.

Appendix A.2. The Automorphism Group of the Cell

The cuboctahedron graph is the line graph of the graph of the cube, $\Gamma_{12} \cong \mathcal{L}(Q_3)$, with the twelve channels corresponding to the twelve edges of the cube. By Whitney's theorem the line-graph construction preserves the automorphism group for connected

graphs on more than four vertices, and the automorphism group of the three-cube is the hyperoctahedral group of order forty-eight. Hence,

$$\Gamma_{12} \cong L(Q_3), \;\; \mathrm{Aut}(\Gamma_{12}) \cong \mathrm{Aut}(Q_3) \cong O_h, \;\; |O_h| = 48. \tag{A4}$$

Two consequences are needed below, and both were verified by exhaustive enumeration rather than assumed. First, the graph admits no automorphism beyond those realised geometrically by the point group: the enumeration returns exactly forty-eight, so the combinatorial symmetry of the cell and its geometric symmetry coincide, and no hidden symmetry of the channel structure has been overlooked. Second, the action is transitive on the twelve vertices, so the twelve channels form a single orbit. This settles the first of the two questions of Section 2.6 in the affirmative.

Appendix A.3. Decomposition of the Bond Space

The permutation character of V_{12} is the number of nearest-neighbour vectors left fixed by each group element. Over the ten classes of O_h it reads

$$\chi_{12} = (12,\ 0,\ 2,\ 0,\ 0,\ 0,\ 0,\ 0,\ 4,\ 2). \tag{A5}$$

The permutation character is resolved over the conjugacy classes of O_h as shown in Table A1.

Table A1. Permutation character of the twelve-channel bond space over the classes of the octahedral group *.

Class	Class Size	Vectors Fixed
E	1	12
$8C_3$	8	0
$6C_2{}'$	6	2
$6C_4$	6	0
$3C_2$	3	0
i	1	0
$6S_4$	6	0
$8S_6$	8	0
$3\sigma h$	3	4
$6\sigma d$	6	2

* Only the two-fold axes lying along ⟨110⟩, the mirror planes normal to the cube axes and the diagonal mirror planes fix any nearest-neighbour vector; every other class fixes none.

Projecting with the orthogonality relation,

$$n_\alpha = \frac{1}{48}\sum_K |K|\,\chi_{12}(K)\chi_\alpha(K), \tag{A6}$$

gives multiplicities that are all zero or one:

$$V_{12} \cong A_{1g} \oplus E_g \oplus T_{2g} \oplus T_{1u} \oplus T_{2u}. \tag{A7}$$

The decomposition is multiplicity-free. Its immediate consequence is that every irreducible subspace is an eigenspace of any O_h-invariant operator on the channel space, and in particular of the Laplacian. Direct diagonalisation gives

$$\mathrm{spec}(A) = \left\{4,\ 2^3,\ 0^3,\ (-2)^5\right\}, \;\; \mathrm{spec}(L) = \left\{0,\ 2^3,\ 4^3,\ 6^5\right\} \tag{A8}$$

with the assignment shown in Table A2, which identifies each Laplacian eigenspace with its irreducible O_h content.

Table A2. Spectrum of the combinatorial Laplacian of the cuboctahedron graph and its decomposition *.

Laplacian Eigenvalue	Multiplicity	Irreducible Content
0	1	A_1g (uniform channel occupation)
2	3	T_1u (vector channel; carries the momentum coupling)
4	3	T_2g
6	5	$Eg \oplus T_2u$

* The four eigenvalues separate all irreducible components except Eg and T_2u, which share the eigenvalue six. That degeneracy is not imposed by the point group: it is a property of the equal-weight adjacency of the cell, and any invariant coupling that weights the two components differently will lift it.

The eigenvector of the vector representation is the projection of the propagation direction on the channels, $c_i = \hat{\mathbf{n}} \cdot \mathbf{e}_i$, for which the neighbour sum closes on the vector itself:

$$c_i = \hat{\mathbf{n}} \cdot \mathbf{e}_i, \quad \sum_{j \sim i} c_j = 2c_i, \qquad \text{i.e.,} \quad A\mathbf{c} = 2\mathbf{c}. \tag{A9}$$

Appendix A.4. Why the Leading Dispersion Term Is Isotropic and the Anisotropy Unique

Section 3.18 states that the space of spherical harmonics of degree two contains no invariant representation of O_h, whereas the degree-four sector contains exactly one. The same character projection used above establishes it:

$$\mathcal{D}^{(2)}|_{O_h} \cong E_g \oplus T_{2g}, \quad \mathcal{D}^{(4)}|_{O_h} \cong A_{1g} \oplus E_g \oplus T_{1g} \oplus T_{2g}. \tag{A10}$$

The space of harmonics of degree two contains no A_{1g} component. There is therefore no O_h-invariant quadratic form on the wave vector other than a multiple of δ_{ij}, and the leading term of the dispersion relation is exactly isotropic for any invariant action on the cell. This is the group-theoretic content of the first row of Table 4: the dimension-four operator $c^{ij}F_{i\mu}F^{\mu}_{j}$ is forced to be proportional to the identity and is reabsorbed into the normalisation of c.

The space of harmonics of degree four contains exactly one invariant component. The leading anisotropy is therefore carried by a single cubic harmonic, which is K_4 of Equation (65), and its angular shape admits no free parameter. The expansion obtained by direct computation in Equation (66),

$$\omega^2 = c^2k^2\left[1 - a^2k^2\left(\frac{1}{40} - \frac{K_4(\hat{n})}{96}\right)\right], \tag{A11}$$

thus has an angular structure fixed by symmetry alone. Only the amplitude—the ratio of the coefficient of K_4 to the isotropic part, equal to 5/12 for the face-centred cubic cell—depends on the dynamics. This is the precise sense of the claim in Section 3.16 that the octahedral pattern is a shape rather than a number, and it is what makes the result of Section 3.13 a consequence of symmetry rather than a feature of the particular expansion carried out there.

The same projection applied at degree six returns exactly one invariant component as well, so the next admissible anisotropic harmonic is likewise unique. The angular structure of the correction is therefore rigid at every order at which it is presently computed, and the freedom left to the dynamics is confined to the coefficients.

Appendix A.5. How Much Freedom the Symmetry Leaves

Because the decomposition is multiplicity-free and every irreducible representation of O_h is of real type, the algebra of invariant operators on the channel space has dimension equal to the number of components,

$$\dim \mathrm{End}_{O_h}(V_{12}) = \sum_{\alpha} n_\alpha^2 = 5, \tag{A12}$$

so the point group admits a five-parameter family of invariant real symmetric couplings on V_{12}. One of those parameters is exhausted by the requirement that the coupling be diagonal in the channels: transitivity then forces equal weights, and the lattice action of Equation (62) is the unique such form up to the normalisation N. The remaining four correspond to couplings between distinct channels—cross-bond terms that Equation (62) does not contain. Their exclusion is a dynamical assumption of the construction, not a consequence of the symmetry, and it belongs alongside the other assumptions recorded in Section 2.1. Such cross-bond couplings preserve the permitted O_h-invariant angular structure but may change its coefficients, including the isotropic part and the sign of the quartic correction. Equation (71) therefore applies only to the diagonal equal-weight nearest-neighbour action of Equation (62).

Appendix A.6. Comparison with the Body-Centred and Simple Cubic Cells

The same computation for the eight nearest neighbours of the body-centred cubic lattice and the six of the simple cubic lattice gives

$$V_8 \cong A_{1g} \oplus T_{2g} \oplus A_{2u} \oplus T_{1u}, \quad V_6 \cong A_{1g} \oplus E_g \oplus T_{1u}. \tag{A13}$$

The resulting bond-space decompositions for the three cubic lattices are compared in Table A3.

Table A3. Bond-space decompositions of the three cubic cells *.

Cell	Channels	Decomposition of the Bond Space
Face-centred cubic	12	$A_1g \oplus Eg \oplus T_2g \oplus T_1u \oplus T_2u$
Body-centred cubic	8	$A_1g \oplus T_2g \oplus A_2u \oplus T_1u$
Simple cubic	6	$A_1g \oplus Eg \oplus T_1u$

* The three are pairwise distinct, so the channel space identifies the cell unambiguously; but all three carry the same point group, so the argument of Appendix A.4 applies verbatim to each of them.

The three decompositions are pairwise distinct, and the twelve-channel space is the only one of the three containing T_{2u}. The bond space therefore identifies the cell. It does not, however, distinguish it in any way that reaches an observable, because all three cells carry the octahedral point group: the argument of Appendix A.4 holds for each, the leading term is isotropic in each, and the leading anisotropy is the same harmonic K_4 in each. The three differ only in the coefficient multiplying it—5/12, 10/9 and −5/3 respectively, the change in sign in the simple cubic case being the inverted angular ordering noted after Equation (72)—which is exactly the result already obtained there by direct computation, and correctly described there as the fingerprint of the lattice.

Appendix A.7. What the Computation Settles

The question of Section 2.6 was posed so as to admit a definite negative answer, and that is the answer it receives. The twelve channels are a single orbit; the bond space decomposes multiplicity-free; the decomposition is distinctive among the cubic cells. However, nothing in it privileges twelve. No property of the decomposition is violated by the eight-channel or

the six-channel cell, and no observable consequence separates them beyond the amplitude already computed in Equation (72). The identification of the maximally symmetric QSR lattice with the Isotropic Vector Matrix is therefore not recovered as a theorem, and coordination twelve remains a defining assumption of the framework rather than a derived result. What the computation contributes is that this negation is now demonstrated rather than merely conceded.

The result does bear on the tension recorded in Section 2.6 between the assumed twelve-fold coordination and the unitarity conditions on the tick operator, and it bears on it in a way worth stating because it narrows the problem. The vector representation T_{1u} occurs exactly once in the twelve-channel space and exactly once in the eight-channel space. The channel that carries the momentum coupling of a two-component walk therefore exists, and exists uniquely, in both. Whatever obstruction favours eight-fold coordination is thus invisible at the level of the permutation representation of the cell: it cannot follow from the symmetry content of the channel space, and must come instead from the unitarity conditions imposed on the coin operator. Settling it requires the tick operator restricted to the T_{1u} channel, which is not constructed here and remains among the priorities of Section 4.6.

Appendix A.8. Numerical Verification

Every statement above has been checked by exhaustive computation on the explicit twelve-, eight- and six-vector sets: the four-regularity and edge count of the channel graph, the order of its automorphism group and the transitivity of the action, the permutation character class by class, the multiplicities of Equation (A7), the Laplacian spectrum and its assignment in Table A2, the eigenvalue relation of Equation (A9), and the restrictions of the degree-two, degree-four and degree-six harmonics. The group was realised as the forty-eight signed permutation matrices in three dimensions, and the character table was constructed and verified orthonormal before use. The isomorphism with the line graph of the cube and the equality of the two automorphism group orders were checked independently by graph isomorphism.

References

1. Bombelli, L.; Lee, J.; Meyer, D.; Sorkin, R.D. Space-time as a causal set. *Phys. Rev. Lett.* **1987**, *59*, 521–524. [CrossRef] [PubMed]
2. Rovelli, C. *Quantum Gravity*; Cambridge University Press: Cambridge, UK, 2004.
3. Ambjørn, J.; Jurkiewicz, J.; Loll, R. Reconstructing the universe. *Phys. Rev. D* **2005**, *72*, 064014. [CrossRef]
4. Dowker, F. Causal sets and the deep structure of spacetime. In *100 Years of Relativity*; World Scientific: Singapore, 2005.
5. Surya, S. The causal set approach to quantum gravity. *Living Rev. Relativ.* **2019**, *22*, 5. [CrossRef]
6. Ambjørn, J.; Görlich, A.; Jurkiewicz, J.; Loll, R. Nonperturbative quantum gravity. *Phys. Rep.* **2012**, *519*, 127–210. [CrossRef]
7. Planck, M. Zur Theorie des Gesetzes der Energieverteilung im Normalspektrum. *Verh. Dtsch. Phys. Ges.* **1900**, *2*, 237–245.
8. Penrose, R. *The Road to Reality*; Jonathan Cape: London, UK, 2004.
9. Hehl, F.W.; von der Heyde, P.; Kerlick, G.D.; Nester, J.M. General relativity with spin and torsion: Foundations and prospects. *Rev. Mod. Phys.* **1976**, *48*, 393–416. [CrossRef]
10. Debever, R. (Ed.) *Elie Cartan–Albert Einstein: Letters on Absolute Parallelism, 1929–1932*; Princeton University Press: Princeton, NJ, USA, 1979.
11. Shapiro, I.L. Physical aspects of the space-time torsion. *Phys. Rep.* **2002**, *357*, 113–213. [CrossRef]
12. Rideout, D.P.; Sorkin, R.D. Classical sequential growth dynamics for causal sets. *Phys. Rev. D* **2000**, *61*, 024002. [CrossRef]
13. Cortês, M.; Smolin, L. Quantum Energetic causal sets. *Phys. Rev. D* **2014**, *90*, 044035. [CrossRef]
14. Ferkiss, V.; Fuller, R.B.; Applewhite, E.J. *Synergetics: Explorations in the Geometry of Thinking*; Macmillan: New York, NY, USA, 1975.
15. Fuller, R.B. *Synergetics 2*; Macmillan: New York, NY, USA, 1979.
16. Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.* **2020**, *641*, A6.
17. Bekenstein, J.D. Black holes and entropy. *Phys. Rev. D* **1973**, *7*, 2333–2346. [CrossRef]
18. Verlinde, E. On the origin of gravity and the laws of Newton. *J. High Energy Phys.* **2011**, *2011*, 29. [CrossRef]
19. Strominger, A. The dS/CFT correspondence. *J. High Energy Phys.* **2001**, *10*, 034. [CrossRef]

20. Einstein, A. On a heuristic point of view concerning the production and transformation of light. *Ann. Phys.* **1905**, *17*, 132–148. [CrossRef]
21. Einstein, A. Does the inertia of a body depend upon its energy content? *Ann. Phys.* **1905**, *18*, 639–641. [CrossRef]
22. de Broglie, L. Recherches sur la Théorie des Quanta. Doctoral Thesis, Université de Paris, Paris, France, 1924.
23. Schrödinger, E. Quantisierung als Eigenwertproblem. *Ann. Phys.* **1926**, *79*, 361–376. [CrossRef]
24. Heisenberg, W. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Z. Phys.* **1927**, *43*, 172–198. [CrossRef]
25. Feynman, R.P. *The Character of Physical Law*; MIT Press: Cambridge, MA, USA, 1965.
26. Wilson, K.G. Confinement of quarks. *Phys. Rev. D* **1974**, *10*, 2445–2459. [CrossRef]
27. Abdo, A.A.; Ackermann, M.; Ajello, M.; Asano, K.; Atwood, W.B.; Axelsson, M.; Baldini, L.; Ballet, J.; Barbiellini, G.; Baring, M.G.; et al. A limit on the variation of the speed of light arising from quantum gravity effects. *Nature* **2009**, *462*, 331–334. [CrossRef] [PubMed]
28. Vasileiou, V.; Jacholkowska, A.; Piron, F.; Bolmont, J.; Couturier, C.; Granot, J.; Stecker, F.W.; Cohen-Tanugi, J.; Longo, F. Constraints on Lorentz invariance violation from Fermi-LAT observations of gamma-ray bursts. *Phys. Rev. D* **2013**, *87*, 122001. [CrossRef]
29. Ackermann, M.; Ajello, M.; Asano, K.; Atwood, W.B.; Axelsson, M.; Baldini, L.; Ballet, J.; Barbiellini, G.; Baring, M.G.; Bastieri, D.; et al. Fermi-LAT observations of the gamma-ray burst GRB 130427A. *Science* **2014**, *343*, 42–47. [CrossRef] [PubMed]
30. Abramowski, A.; Acero, F.; Aharonian, F.; Akhperjanian, A.; Anton, G.; Barnacka, A.; de Almeida, U.B.; Bazer-Bachi, A.; Becherini, Y.; Becker, J.; et al. Search for Lorentz invariance breaking with a likelihood fit of the PKS 2155-304 flare data taken on MJD 53944. *Astropart. Phys.* **2011**, *34*, 738–747. [CrossRef]
31. CTA Consortium. *Science with the Cherenkov Telescope Array*; World Scientific: Singapore, 2019. [CrossRef]
32. Cao, Z.; Aharonian, F.; Axikegu; Bai, Y.X.; Bao, Y.W.; Bastieri, D.; Bi, X.J.; Bi, Y.J.; Bian, W.; Bukevich, A.V.; et al. Stringent Tests of Lorentz Invariance Violation from LHAASO Observations of GRB 221009A. *Phys. Rev. Lett.* **2024**, *133*, 071501. [CrossRef] [PubMed]
33. Myers, R.C.; Pospelov, M. Ultraviolet modifications of dispersion relations in effective field theory. *Phys. Rev. Lett.* **2003**, *90*, 211601. [CrossRef] [PubMed]
34. Toma, K.; Mukohyama, S.; Yonetoku, D.; Murakami, T.; Gunji, S.; Mihara, T.; Morihara, Y.; Sakashita, T.; Takahashi, T.; Wakashima, Y.; et al. Strict limit on CPT violation from polarization of gamma-ray bursts. *Phys. Rev. Lett.* **2012**, *109*, 241104. [CrossRef] [PubMed]
35. Laurent, P.; Götz, D.; Binétruy, P.; Covino, S.; Fernandez-Soto, A. Constraints on Lorentz invariance violation using INTEGRAL/IBIS observations of GRB 041219A. *Phys. Rev. D* **2011**, *83*, 121301. [CrossRef]
36. Kostelecký, V.A.; Mewes, M. Electrodynamics with Lorentz-violating operators of arbitrary dimension. *Phys. Rev. D* **2009**, *80*, 015020. [CrossRef]
37. Burnett, J.H.; Levine, Z.H.; Shirley, E.L. Intrinsic birefringence in calcium fluoride and barium fluoride. *Phys. Rev. B* **2001**, *64*, 241102(R).
38. Kostelecký, V.A.; Russell, N. Data tables for Lorentz and CPT violation. *Rev. Mod. Phys.* **2011**, *83*, 11–31. [CrossRef]
39. Collins, J.; Perez, A.; Sudarsky, D.; Urrutia, L.; Vucetich, H. Lorentz invariance and quantum gravity: An additional fine-tuning problem? *Phys. Rev. Lett.* **2004**, *93*, 191301. [CrossRef] [PubMed]
40. Bombelli, L.; Henson, J.; Sorkin, R.D. Discreteness without symmetry breaking: A theorem. *Mod. Phys. Lett. A* **2009**, *24*, 2579–2587. [CrossRef]
41. Lange, R.; Huntemann, N.; Rahm, J.M.; Sanner, C.; Shao, H.; Lipphardt, B.; Tamm, C.; Weyers, S.; Peik, E. Improved limits for violations of local position invariance from atomic clock comparisons. *Phys. Rev. Lett.* **2021**, *126*, 011102. [CrossRef] [PubMed]
42. Damour, T.; Dyson, F. The Oklo bound on the time variation of the fine-structure constant revisited. *Nucl. Phys. B* **1996**, *480*, 37–54. [CrossRef]
43. Murphy, M.T.; Molaro, P.; Leite, A.C.O.; Cupani, G.; Cristiani, S.; D'Odorico, V.; Santos, R.G.; Martins, C.J.A.P.; Milaković, D.; Nunes, N.J.; et al. Fundamental physics with ESPRESSO: Precise limit on variations in the fine-structure constant towards the bright quasar HE 0515−4414. *Astron. Astrophys.* **2022**, *658*, A123. [CrossRef]
44. Bialas, P.; Burda, Z.; Krzywicki, A.; Petersson, B. Focusing on the fixed point of 4D simplicial gravity. *Nucl. Phys. B* **1996**, *472*, 293–308. [CrossRef]
45. de Bakker, B.V. Further evidence that the transition of 4D dynamical triangulation is first order. *Phys. Lett. B* **1996**, *389*, 238–242. [CrossRef]
46. Ambjørn, J.; Jurkiewicz, J.; Loll, R. Emergence of a 4D world from causal quantum gravity. *Phys. Rev. Lett.* **2004**, *93*, 131301. [CrossRef] [PubMed]
47. Ambjørn, J.; Görlich, A.; Jurkiewicz, J.; Loll, R. Planckian birth of the quantum de Sitter universe. *Phys. Rev. Lett.* **2008**, *100*, 091304. [CrossRef] [PubMed]

48. Ambjørn, J.; Jordan, S.; Jurkiewicz, J.; Loll, R. Second-order phase transition in causal dynamical triangulations. *Phys. Rev. Lett.* **2011**, *107*, 211303. [CrossRef] [PubMed]
49. Loll, R. Quantum gravity from causal dynamical triangulations: A review. *Class. Quantum Gravity* **2020**, *37*, 013002. [CrossRef]
50. Klitgaard, N.; Loll, R. Introducing quantum Ricci curvature. *Phys. Rev. D* **2018**, *97*, 046008. [CrossRef]
51. Loomis, S.P.; Carlip, S. Suppression of non-manifold-like sets in the causal set path integral. *Class. Quantum Grav.* **2018**, *35*, 024002. [CrossRef]
52. Regge, T. General relativity without coordinates. *Nuovo Cim.* **1961**, *19*, 558–571. [CrossRef]
53. Hořava, P. Quantum gravity at a Lifshitz point. *Phys. Rev. D* **2009**, *79*, 084008. [CrossRef]
54. Ambjørn, J.; Görlich, A.; Jordan, S.; Jurkiewicz, J.; Loll, R. CDT meets Hořava–Lifshitz gravity. *Phys. Lett. B* **2010**, *690*, 413–419. [CrossRef]
55. Conway, J.H.; Sloane, N.J.A. *Sphere Packings, Lattices and Groups*, 3rd ed.; Springer: New York, NY, USA, 1999. [CrossRef]
56. Edwards, R.G.; Joó, B.; Lin, H.-W. Tuning for three flavors of anisotropic clover fermions with stout-link smearing. *Phys. Rev. D* **2008**, *78*, 054501. [CrossRef]
57. Borel, A. Density properties for certain subgroups of semi-simple groups without compact components. *Ann. Math.* **1960**, *72*, 179–188. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.