On the Absence of Automatic Linear Lorentz-Violating Dispersion from discreteness Alone

A note on emergent-metric discrete spacetime models

(C)July 2026 Sergej Materov*

Abstract

We establish a structural criterion determining when a discrete spacetime model does or does not produce a first-order, energy-dependent (linear-in-E) Lorentz-violating photon dispersion. The criterion consists of four conditions defining the emergent-metric class: absence of a primitive metric, absence of a canonical background embedding, observer-relative local observables, and a decorrelation property of the reconstruction procedure. We prove two theorems. First, every member of this class satisfying all four conditions is structurally protected from linear Lorentz-violating dispersion: the accumulated phase grows as the square root of the number of traversed elements, suppressing any first-order signal. Second, for the sub-class of models whose generating point process is Poincaré-invariant in law – which includes Causal Set Theory – this protection is proven to be preserved under Lorentz boosts, so no inertial observer sees a coherent linear dispersion signal regardless of velocity; for the remaining members of the emergent-metric class, boost invariance is stated as an explicit, falsifiable conjecture rather than derived from (EM1)–(EM3) alone. Models outside the class - those assuming a fixed, rigid embedding with a preferred orientation are not structurally protected against coherent linear dispersion, which may then arise depending on additional model-specific assumptions. The result is substrate-independent. Its primary value is not the explanation of any particular observation, but the identification of the precise structural boundary between discrete models that generically predict linear Lorentz violation and those that do not. This paper is the first in a series establishing a general framework for discrete quantum gravity theories with emergent metric structure.

1. Motivation

The central question of this paper is structural: under what precise conditions does a discrete spacetime model produce, or fail to produce, a first-order energy-dependent photon dispersion? This question is often conflated with a phenomenological one - whether current GRB observations rule out discrete spacetime but the two are logically distinct. As Amelino-Camelia and others have noted, theorists working in discrete quantum gravity did not generally expect a strong implication from discreteness to in-vacuo dispersion; such an expectation appears mainly in phenomenological discussions that treat discreteness as equivalent to a fixed lattice. The present paper makes the structural distinction precise.

Two properties of a discrete model are logically independent: discreteness of the underlying substrate, and the existence of a fixed, rigid embedding of that substrate into a metric space with a privileged orientation. The latter, not the former, is what produces a generic linear dispersion signal. Causal set theory illustrates this at the level of the substrate: a Lorentz-invariant Poisson sprinkling avoids a preferred frame by construction. The present paper generalises the underlying mechanism to an explicitly defined class of models - the emergent-metric class (EM1)–(EM4) and establishes the precise structural boundary: models inside the class do not generically produce linear dispersion; models outside it do. The GRB phenomenology is one application of this boundary, but the boundary is the result.

*Independed Researcher E-mail: sergejmaterov2@gmail.com

2. Setup and Definition of the Emergent-Metric Class

Let *G* be a discrete substrate (a graph, causal set, or similar combinatorial structure) without a primitive metric. Physical distance *D* and duration are assumed to be reconstructed quantities, obtained from *G* via some statistical or dynamical procedure *R* (e.g. diffusion processes, correlation decay, or information-theoretic measures on *G*), rather than read off a fixed embedding chosen in advance.
Propagation of a massless excitation between two reconstructed points is modeled as traversal of *N* local configurations of *G*, each contributing a phase $\delta\varphi_i$ that depends on the local orientation of that configuration relative to some reference direction.

Definition (Emergent-Metric Class). A discrete spacetime model (G, R) belongs to the *emergent-metric class* if it satisfies all four of the following conditions:

(EM1) Reconstructed metric: the physical metric: distances, durations, and angles - is not primitively assigned to *G* but is derived entirely from *G* via the procedure *R*.

(EM2) No canonical embedding: there exists no observer-independent, canonical embedding of *G* into a fixed background metric space. The reconstructed manifold *M = R(G)* is the only spacetime structure available to observers.

(EM3) Relative local observables: all local physical observables are defined relative to the reconstructed coordinates of *M*, not relative to any coordinate system on *G* itself. In particular, the orientation of a local configuration of *G* is defined only via *R*, not intrinsically.

(EM4) Decorrelation property: the reconstruction procedure *R* induces, at the epistemic level of any observer or propagating excitation, a distribution of local phase contributions $\delta\varphi_i$ satisfying

$$\langle \delta\varphi_i \delta\varphi_j \rangle \approx\ 0,\ \ i\ \neq\ j,$$

i.e. the orientation-dependent phase contributions along any reconstructed path are uncorrelated at long range. This is an equidistribution property of *R*, not an assumption about the dynamics of *G*. Conditions (EM1)–(EM3) are necessary for the emergent-metric class; condition (EM4) is the additional structural requirement under which the main theorems below hold. Models satisfying only (EM1)–(EM3) but not (EM4) belong to the class but do not necessarily satisfy the theorems.

A clarification is in order before stating the theorem, since the decorrelation condition used below is easy to misread as an appeal to chaotic or random underlying dynamics. It is not. Decorrelation of phase contributions, $\langle \delta\varphi_i \delta\varphi_j \rangle \approx\ 0$, is a statement about equidistribution - the statistical spread of orientations sampled along a path - not about sensitivity to initial conditions. A fully deterministic, non-chaotic, unitary dynamics can still produce a sequence of orientations that is equidistributed and effectively uncorrelated in the relevant statistical sense; the digits of an irrational rotation on a torus are a standard example of a deterministic, non-chaotic sequence with exactly this property. The decorrelation assumed here is therefore compatible with - and, in a fully unitary microscopic theory, should be understood as arising from an epistemic coarse-graining performed by any observer or propagating excitation that cannot resolve the full microscopic configuration of G, rather than from any genuine, ontic chaos in the underlying dynamics. This distinction matters: it keeps the present argument consistent with microscopic theories that are deterministic and non-chaotic at the ontic level, while still producing apparent randomness at the coarse-grained, epistemic level relevant to propagation.

Main Result

The two theorems below (Theorem 1 and the Boost-Invariance Lemma of Section 6) together establish the following combined result, which is the central claim of this paper:

Main Theorem (Meta Theorem). For every member of the emergent-metric class satisfying conditions (EM1)–(EM4), linear Lorentz-violating photon dispersion ($\Phi_N \propto N$, equivalently a dispersive term linear in $\frac{E}{E_{Pl}}$) is absent in every inertial frame. No coherent dispersion signal of the first order accumulates along any reconstructed path for any observer, regardless of their velocity.

This result has three components: (i) the decorrelation condition (EM4) prevents coherent accumulation of phase in any fixed frame (Theorem 1 below); (ii) for the Poincaré-invariant-in-law sub-class – which includes Causal Set Theory – the decorrelation condition is proven to be preserved under Lorentz boosts (Lemma D, §6.1); for the remaining members of the class this is stated as Conjecture 2′ (§6.4), since conditions (EM1)–(EM3) alone do not by themselves establish it; (iii) therefore no inertial observer, regardless of velocity, sees a coherent first-order dispersion signal. The class defined by (EM1)-(EM4) is thus structurally protected from linear LIV, not by fine-tuning or model-specific cancellation, but by the geometry of the reconstruction procedure itself.

Remark on prior work. The physical intuition that non-systematic Lorentz violation averages rather than accumulates is present in the literature, notably in Basu & Mattingly (2005) and in the causal set context (Dowker, Henson & Sorkin 2004). The present paper adds three things not found there: (a) the formal definition of the emergent-metric class (EM1)–(EM4) as the precise domain of validity; (b) the proof, for the Poincaré-invariant-in-law sub-class including Causal Set Theory, that condition (EM4) is preserved under boosts, together with an explicit conjecture for the remaining members of the class (Section 6) – Basu & Mattingly do not address this question at all, working instead in a fixed preferred frame; and (c) the model-independent observational prediction of $\sqrt{D}$-scaling (Section 7), which is a consequence of the Main Theorem rather than a separate phenomenological assumption.

Theorem 1: Absence of Automatic Linear Dispersion from Discreteness

A first-order Lorentz-violating dispersion relation, i.e. an accumulated phase $\Phi_N = \sum_{n=1}^{N} \delta\varphi_i$ growing linearly, $\Phi_N \propto N$, equivalently a dispersive term linear in $\frac{E}{E_{Pl}}$ - is ***not*** a generic consequence of the discreteness of *G* alone. It requires an additional, independent structural assumption: that the local phase contributions $\delta\varphi_i$ are coherently aligned along the propagation direction, i.e. $\langle\delta\varphi_i\rangle \neq 0$ with a consistent sign set by a globally preferred frame.

Proof sketch. If *R* does not single out a preferred frame - i.e. there is no canonical, fixed embedding of *G* into a metric space, consistent with the metric being emergent - then the orientation of local configurations sampled along a reconstructed path is not fixed relative to the propagation direction. Suppose, in addition, the local phase contributions are uncorrelated at long range:

$$\left\langle\delta\varphi_i\delta\varphi_j\right\rangle \approx 0, \ \ i \neq j.$$

This decorrelation condition is itself a separate, checkable property of the reconstruction procedure *R* - specifically, an equidistribution property of orientations under *R*, recovered at the epistemic, coarse-grained level rather than presupposed as ontic randomness or chaos (see the remark above) - and not an assumption about discreteness per se. Under this condition, the variance of the accumulated phase is additive,

$$Var(\Phi_N) = N \cdot Var(\delta\varphi), \ \ so \ \ \sigma_\Phi \sim \sqrt{N},$$

rather than $\sigma_\Phi \sim N$. A linear-in-N (coherent) signal therefore requires either (a) a fixed background embedding, which breaks the emergent-metric assumption, or (b) explicit long-range phase correlation $\left\langle\delta\varphi_i\delta\varphi_j\right\rangle \neq 0$ surviving reconstruction. Neither follows from discreteness of *G* by itself. ■

3. Corollary and Scope

Corollary. Observed linear Lorentz-violating dispersion (or its absence) is informative about discrete models *only* to the extent that those models additionally assume a fixed, rigid embedding with a privileged orientation. Models with emergent, reconstructed, and dynamically reconfiguring metrics are not automatically constrained by GRB linear-dispersion bounds in the way a static crystalline lattice would be; the bound constrains the coherence assumption, not discreteness as such.

Scope. The argument is purely structural: it uses no feature specific to any particular discrete-gravity proposal beyond the absence of a fixed embedding and the decorrelation property $\langle \delta\varphi_i \delta\varphi_j \rangle \approx 0$. It therefore applies to any discrete substrate with an emergent, reconstructed metric satisfying this property, including (but not limited to) causal set theory and related emergent-geometry approaches to quantum gravity. The theorem is accordingly a *consistency result* protecting an entire class of theories from a specific objection - it does not, by itself, provide evidence for any individual member of that class, and does not by itself constitute a falsifiable prediction. Discriminating between specific models in this class requires additional, model-specific structure beyond what is assumed here.

Conditional result. The theorem establishes that linear dispersion is not *forced* by discreteness; it does not establish that it is *absent* in any specific model. For a given discrete substrate and reconstruction procedure *R*, verifying the decorrelation property $\langle \delta\varphi_i \delta\varphi_j \rangle \approx 0$ explicitly, and ruling out any residual preferred frame surviving reconstruction, remains a separate, model-specific task.

3b. Relation to Prior Work and the CLT Objection

Relation to the central limit theorem. The mathematical mechanism in Theorem 1 is additive variance for uncorrelated variables: $\sigma_\Phi \sim \sqrt{N}$. This is the CLT. The contribution of the present paper is not this mechanism but its domain of validity in quantum gravity. Three things go beyond the CLT: (i) conditions (EM1)–(EM4) constitute a structural criterion identifying which discrete quantum gravity theories satisfy independence - this requires a physical argument about background-independence as a necessary condition; (ii) Theorem 2 proves, for the Poincaré-invariant-in-law sub-class including Causal Set Theory, that (EM4) is preserved under Lorentz boosts – a separate argument required by the non-compactness of the boost group – and identifies the precise, checkable condition (Conjecture 2′) under which this extends to the rest of the emergent-metric class; (iii) the boundary is sharp - models violating (EM2) produce coherent accumulation regardless of other assumptions. The CLT is a tool; the result is the identification of when it applies to quantum gravity, and that it applies in all inertial frames.

We now discuss the six bodies of prior work most directly relevant to the present results, identified by the pre-reviewer as the principal potential competitors.

Bombelli, Henson & Sorkin (2009). Bombelli et al. prove that Lorentz-invariant Poisson sprinklings cannot intrinsically define a preferred timelike direction or preferred inertial frame. The present work starts from a different premise: rather than proving the absence of preferred directions for one specific discrete construction, it identifies general structural conditions (EM1–EM4) under which coherent linear dispersion is impossible for an entire class of emergent-metric theories.

Dowker, Henson & Sorkin (2004). This paper shows that for causal sets with a Lorentz-invariant Poisson sprinkling, discreteness does not produce systematic Lorentz violation. This is arguably the closest predecessor to Theorem 1. The Poisson sprinkling satisfies (EM1)–(EM3) by construction, and the statistical independence provides a natural mechanism supporting (EM4). The present paper generalises this result in two ways: (a) to the entire emergent-metric class (EM1)–(EM4), not only causal sets; and (b) by proving that the decorrelation condition is preserved under boosts (Theorem 2), which is not addressed by Dowker et al. To the best of our knowledge, no published theorem formulates an equivalent structural criterion in terms of conditions (EM1)–(EM4) together with boost invariance.

Sorkin and related causal set phenomenology. Sorkin's programme establishes causal sets as a background-independent discrete substrate and discusses the emergence of Lorentz invariance in the continuum limit. The present paper's conditions (EM1)–(EM3) can be viewed as an abstraction of what causal set theory achieves by construction: no fixed embedding, no preferred frame, reconstructed metric. The present contribution is not another result about causal sets themselves, but the identification of a substrate-independent structural class characterised by (EM1)–(EM4). The present contribution is not another result about causal sets themselves, but the identification of a substrate-independent structural class characterised by (EM1)–(EM4), together with the proof that these conditions exclude coherent linear LIV and remain valid under Lorentz boosts for the Poincaré-invariant-in-law sub-class, with boost invariance for the wider emergent-metric class stated as an explicit conjecture (Section 6). To the best of our knowledge, no equivalent structural theorem has been published for this sub-class.

Amelino-Camelia (DSR and quantum-spacetime phenomenology). Amelino-Camelia's doubly special relativity (DSR) programme modifies Lorentz transformations to make the Planck length an observer-independent invariant, and his phenomenological framework parametrises LIV effects through deformation coefficients η and n. This is a different approach: DSR modifies the symmetry group, whereas the emergent-metric class preserves it (as an approximate symmetry at accessible scales). The present theorem is not in conflict with DSR but is structurally distinct: we do not deform Lorentz symmetry, we identify conditions under which it is effectively preserved. Amelino-Camelia's phenomenological framework provides no structural criterion distinguishing models with and without coherent accumulation; the emergent-metric class does.

Jacobson, Liberati & Mattingly (2003, 2006). This series of papers derives astrophysical constraints on LIV within the effective field theory (EFT) framework, adding all renormalisable Lorentz-violating operators to the Standard Model and constraining their coefficients from threshold reactions, synchrotron emission, and GRB timing. Their framework explicitly assumes a fixed preferred background frame (the EFT vacuum has a fixed Lorentz-violating tensor), which violates condition (EM2). Their constraints are therefore applicable to models outside the emergent-metric class. The present Corollary makes this explicit: GRB and threshold constraints as derived in the Jacobson–Liberati–Mattingly framework constrain the coherence and embedding assumptions built into the EFT approach, not the emergent-metric class. This does not reduce the value of their constraints; it clarifies their logical scope.

Liberati & Maccione (2011, 2014). These papers extend the EFT LIV constraint programme to higher-order operators and consider dissipation alongside dispersion. As with Jacobson et al., the framework assumes a fixed preferred background, violating (EM2). The present theorem does not conflict with their results; it bounds their domain of applicability. One potential point of genuine tension: Liberati & Maccione discuss "condensate" or emergent-gravity scenarios and argue that even in such scenarios LIV operators may be generated radiatively. If correct, this would imply that some models satisfying (EM1)–(EM3) fail to satisfy (EM4) due to radiatively generated correlations. This is a specific, checkable prediction and represents the most technically serious potential challenge to the present framework. Addressing it requires showing that radiative corrections in models satisfying (EM1)–(EM3) do not generate long-range phase correlations — a model-specific computation that the present theorem correctly identifies as an open task (see Conditional Result in Section 3).

Basu & Mattingly (2005). This paper analyses non-systematic LIV for UHECR protons in a fixed preferred frame (the CMBR rest frame) and shows that coherent accumulation is exponentially unlikely, yielding a weak constraint on velocity fluctuations. The core physical intuition — that non-systematic fluctuations average rather than accumulate - overlaps with Theorem 1. The differences are: (a) Basu & Mattingly work in a fixed preferred frame, violating (EM2); their result is frame-specific and does not address whether the decorrelation survives a boost; (b) they do not provide a structural definition of the class of models for which the averaging holds; (c) they find that particle production in a fluctuating background makes any constraint unreliable, which is a separate physical effect not addressed by the present theorem (whose scope is photon phase accumulation, not particle production). The present paper can be read as providing the structural foundation that Basu & Mattingly's argument implicitly assumes, and extending it to all inertial frames.

4. Relation to Observational Bounds

Current GRB-based bounds constrain the quantum-gravity energy scale associated with linear-in-energy photon dispersion to be at or above the Planck energy. The theorem above clarifies the logical status of these bounds with respect to discrete spacetime models in general: they constrain models that assume coherent phase accumulation relative to a fixed frame, and are silent — to leading structural order — on models in which the metric is emergent and the decorrelation property holds. This does not weaken the observational bounds; it sharpens what class of theoretical models they actually constrain.

5. A Second, Distinct Prediction: Stochastic Spreading vs. Systematic Dispersion

It is important not to conflate two logically distinct observables. Systematic dispersion is a shift in the mean arrival-time difference between photons of different energies, growing as $\langle \Delta t \rangle \propto E$ and tested directly by the GRB timing bounds discussed above. Stochastic wavefront spreading (sometimes called spacetime 'fuzziness') is a different quantity: the random spread of arrival times around the mean, $\sigma(\Delta t)$, which can grow with distance even when the mean shift is exactly zero. This has been tested separately, for instance by modeling photon speeds as normally distributed around c with an energy-dependent standard deviation.

The theorem above bears directly on this second quantity, and the connection is worth making explicit rather than left implicit. The decorrelation condition $\langle \delta\varphi_i \delta\varphi_j \rangle \approx 0$ does not merely suppress the *mean* accumulated phase to zero; it leaves a nonzero *variance*, $\sigma_\Phi \sim \sqrt{N}$, growing with the number of traversed configurations. This residual variance is precisely the stochastic counterpart of wavefront spreading: the theorem predicts not an exactly null result for $\sigma(\Delta t)$, but a specific, suppressed scaling for it, distinct from -and weaker than is the naive coherent estimate that would follow from a fixed embedding.

Two consequences follow. First, the suppression of systematic dispersion established here does not, by itself, establish suppression of stochastic spreading to the same order; the two must be bounded separately, by translating σ_Φ into a time-domain quantity $\sigma(\Delta t)$ with correct dimensional factors and comparing it against the empirical fuzziness bound. Second, because σ_Φ is manifestly nonzero, the theorem does not predict an exact absence of spreading only a strongly suppressed one, consistent with current non-detection but leaving, in principle, a residual, statistically recoverable signal at higher sensitivity. This sharpens the falsifiability of the framework: a future detection of stochastic spreading at a level inconsistent with the $\sqrt{N}$ suppression derived here, but inconsistent also with a null result, would constrain the decorrelation assumption itself.

6. Theorem 2: Boost Invariance of the Decorrelation Condition

The theorem established above suppresses linear Lorentz-violating dispersion in any given inertial frame. A natural and deeper question is whether the decorrelation condition itself – the structural assumption on which the theorem rests – is preserved under Lorentz boosts. If it is not, a boosted observer could in principle see coherent phase accumulation even when a stationary observer does not, restoring a frame-dependent dispersion signal. This section establishes a rigorous answer for a specific, well-defined sub-case of the emergent-metric class – Poincaré-invariant-in-law point processes equipped with a causal-order-equivariant reconstruction procedure, a sub-case that includes Causal Set Theory – and states the general case for the remaining members of Appendix I's group A as an explicit, falsifiable conjecture rather than a claim derived from (EM1)–(EM3) alone.

6.1 Rigorous Lemma: boost invariance for Poincaré-invariant-in-law substrates

Definition (Poincaré-invariant in law). A locally finite point process Ξ on Minkowski space $M^{\{1,d\}}$ is Poincaré-invariant in law if, for every element g of the Poincaré group (Lorentz transformations and translations), the transformed process gΞ has exactly the same probability distribution as Ξ itself.

Fact (standard; Bombelli, Lee, Meyer & Sorkin 1987). A Poisson point process on $M^{\{1,d\}}$ with intensity measure proportional to the Lorentz-invariant spacetime volume form is Poincaré-invariant in law. This holds because a Poisson process's law is completely fixed by its intensity measure, and every Poincaré transformation preserves the Minkowski volume form exactly (unit Jacobian determinant); this is the standard fact underlying the Lorentz-invariance of causal-set sprinklings, established well before and independently of the present paper.

Definition (causal-order-equivariant statistic). A local reconstruction ingredient R*i* is causal-order-equivariant if its value at point *i* is built only from the causal order and proper-time intervals within a bounded neighborhood of *i* - both Lorentz-invariant primitives - so that R*i*(gΞ) is a fixed, deterministic function of gΞ for every g, with no separate reference to any embedding coordinate.

Lemma D (rigorous boost invariance). Let Ξ be Poincaré-invariant in law and let R be causal-order-equivariant. Then for every g in the Poincaré group, the joint law of {R*i*(gΞ)} over any finite index set equals the joint law of {R*i*(Ξ)}. In particular, the mean E[δφ*i*], the covariance ⟨δφ*i*δφ*j*⟩, and the decorrelation property of (EM4) computed in the boosted frame equal their values in the original frame.

Proof. By equivariance, {R*i*(gΞ)} is a fixed measurable functional of gΞ, with no further g-dependence. Since

$$g\Xi =^{d} \Xi \quad \textit{(Poincaré-invariance in law)}$$

any measurable functional of gΞ has, by definition of equality in law, exactly the same distribution as that functional applied to Ξ. Hence $\{R_i(g\Xi)\}$ and $\{R_i(\Xi)\}$ are equal in law. ■

Nothing here presupposes "no preferred frame survives" as an assumption. That conclusion is a consequence of two separately checkable facts - (i) the substrate's law is Poincaré-invariant (a property of the generating process, established independently for Poisson sprinklings), and (ii) R is equivariant (a property of the chosen reconstruction procedure, verified below by direct computation). Causal Set Theory, with its standard Lorentz-invariant Poisson sprinkling (Appendix I, entry A1), satisfies both hypotheses for the natural choices of R used in that literature (link structure, nearest-neighbor intervals, and similar causal-order-built statistics). Theorem 2 a proven result for that sub-class - not merely asserted by general appeal to EM1–EM3.

6.2 Numerical verification

Two properties were checked directly, in 1+1D Minkowski space, using a Poisson sprinkling and the concrete equivariant statistic R_i = the causal-order-nearest predecessor's local velocity $u_i = \Delta x/\Delta t$ (predecessor selected by minimal proper-time interval $\Delta\tau^2 = \Delta t^2 - \Delta x^2$, both Lorentz-invariant quantities):

Check	Method	Result
Predecessor identity preserved under boost	Exact algebraic (2957 points, β=0.6)	100.0% identical

Check	Method	Result
u_{boost} = relativistic velocity addition of u_{orig}	Exact algebraic	max error 4.4×10^{-13} (machine precision)
Naive test, small box, no margin control	KS two-sample, native vs. boosted	D=0.093, p≈0 (spurious - see §6.3)
Corrected test, large box + interior margin	KS two-sample, native vs. boosted	D=0.022, p=0.30
Baseline: native vs. independent native draw	KS two-sample	D=0.019, p=0.47

Table 2. Numerical checks supporting Lemma D (β = 0.6 boost; see §6.3 for the boundary-artifact caveat).

The first two rows confirm equivariance exactly: since causal order and proper time are themselves Lorentz-invariant, the identity of the nearest causal predecessor is preserved bit-for-bit under a boost (verified on 2,957 points), and the transformed local velocity matches the relativistic velocity-addition formula $u' = \frac{u+\beta}{1+u\beta}$ to machine precision - exactly the equivariance Lemma D requires, with no statistical uncertainty at all (see the GitHub repository in the "Reference" section).

6.3 A methodological caveat found in the process: finite-box boundary artifacts

An initial distributional test compared the native (unboosted) empirical distribution of u_i to the distribution obtained by boosting a fresh sample and recomputing u_i in the boosted coordinates. The first attempt, sampling query points from a fixed rectangular box without margin, showed an apparent violation: the boosted-frame mean shifted from ≈0.001 to ≈0.17, with a two-sample Kolmogorov–Smirnov test rejecting equality of distributions (p≈0). The cause was identified and is worth recording explicitly, since it is a generic trap for numerical tests of this kind: a Lorentz boost maps a rectangular sampling box into a sheared parallelogram, so query points near the transformed edges are missing some of the causal predecessors they would have had in a truly homogeneous, unbounded process - a finite-volume censoring effect that mimics a coherent directional bias (an aberration-like artifact) without being one.

Repeating the test with a substantially larger box and restricting query points to an interior region, with margin much larger than both the boost-induced shear and the typical predecessor scale, removed the effect entirely: the corrected KS statistic (D=0.022, p=0.30) is statistically indistinguishable from the native-vs-native baseline (D=0.019, p=0.47) computed from two independent unboosted draws. This confirms Lemma D's prediction directly and should be stated as an explicit methodological requirement for any future numerical test of (EM4) or its boost-invariance in a specific model (e.g. a CDT or spin-foam simulation): finite-region sampling combined with a boost will generically produce a spurious coherent-looking bias unless query points are kept well inside the sampled region, with margin large compared to both the relevant local length scale and the boost-induced shear of the sampling boundary.

For rotations, an equidistribution argument of this kind is automatic: SO(3) is compact, and its invariant Haar measure guarantees $(\Lambda)_{+}\mu = \mu$ exactly, regardless of the specific substrate. For boosts the argument is different in character: no finite invariant measure exists on the non-compact Lorentz boost group, so this route is unavailable in general. Lemma D sidesteps the obstruction for the Poincaré-invariant-in-law sub-class by working directly with equality in law of the point process itself, rather than with an invariant measure on individual configurations – but that route depends on the substrate's generating law already being known to be Poincaré-invariant, which is not established for CDT, LQG/spinfoams, asymptotic safety, or Group Field Theory. Non-compactness of the boost group is sometimes read as evidence that decorrelation must fail at large boosts, with phase-distribution 'tails' coherently aligning along the direction of motion; this inference does not follow. Non-compactness rules out the equidistribution argument used for rotations – it does not, by itself, imply that $\langle\delta\varphi_i\delta\varphi_j\rangle$ acquires a non-zero mean under a boost. The two claims are logically independent, and Lemma D shows that the equidistribution argument survives non-compactness whenever the substrate's law is itself boost-invariant. Whether that premise holds for the remaining members of the emergent-metric class is exactly the open question Conjecture 2′ below makes explicit.

6.4 Conjecture 2′: the general emergent-metric class

For models outside the Poincaré-invariant-in-law point-process sub-class: CDT, LQG/spinfoams, asymptotic safety, Group Field Theory, TQGT*, and any other member of Appendix I's group A beyond causal sets - Lemma D's hypotheses are not currently known to hold, and nothing in conditions (EM1)–(EM3) implies that they do. Absence of a fixed background embedding (EM2) is a statement about the individual theory's kinematics; it does not by itself imply that the theory's ensemble (the path integral, the spinfoam sum, the group-field condensate, etc.), suitably continued to Lorentzian signature, is Poincaré-invariant in law in the sense required by Lemma D. This is an additional, separate, and in general open dynamical question.

Conjecture 2′. For a member of the emergent-metric class satisfying (EM1)–(EM4) in some fixed frame, the decorrelation condition (EM4) continues to hold in every boosted frame, provided the theory's ensemble is Poincaré-invariant in law and its reconstruction procedure R is causal-order-equivariant in the sense of §6.1. Both provisos are open, model-specific questions for every entry in Appendix I marked "EM4 open" (CDT, LQG, asymptotic safety, GFT, TQGT*); establishing them is a strictly harder (it falls outside the scope of the present work and should be carried out by the authors of the theories themselves) requirement than establishing (EM4) itself in a single frame, since (EM4) alone is a single-frame statement and does not automatically transport across observers without the additional law-invariance property identified here.

This sharpens, rather than weakens, the paper's own classification: it identifies exactly what "EM4 open" must be upgraded to (law-invariance of the ensemble under the boost subgroup, not merely single-frame decorrelation) before Theorem 2 can be claimed for any of those models, and it explains why Causal Set Theory alone earns the unqualified "within scope (proven)" status in Appendix I's table: it is the only entry for which both the point-process law-invariance and a natural equivariant R are independently established in the literature.

Corollary (Causal Set Theory). For Causal Set Theory specifically – where the Poisson sprinkling is Poincaré-invariant in law by construction (Bombelli, Lee, Meyer & Sorkin 1987) and natural reconstruction statistics (link structure, nearest-neighbor intervals, and the predecessor-velocity statistic verified numerically in §6.2) are causal-order-equivariant – Lemma D applies directly, and the $\sqrt{N}$ suppression of linear Lorentz-violating dispersion established in Theorem 1 is a boost-invariant, frame-independent result. This is currently the only member of Appendix I's group A for which boost invariance of the decorrelation condition is a proven, rather than conjectured, property.

Falsifiability structure. The content of this section is not that observers built from the same fields as the substrate trivially see no violation – that would be a tautology. For Causal Set Theory, Lemma D gives a specific, checkable claim: it would be falsified by exhibiting a causal-order-equivariant R whose two-point function is zero in one frame but non-zero in a boosted frame, which would falsify equivariance itself rather than the theorem's logic. For the remaining members of the emergent-metric class, Conjecture 2′ is falsifiable in the same sense used throughout this paper: exhibiting a model in Appendix I's group A whose ensemble law fails to be Poincaré-invariant, or whose natural R fails to be causal-order-equivariant, and for which (EM4) nonetheless fails to survive a boost, would refute the conjecture for that model without touching the proven Causal Set Theory case. The framework retains the three independent falsification channels identified in Section 3b: (a) explicit verification that a specific R satisfies (EM4) in a given model – a model-specific computation; (b) whether radiative corrections generate long-range phase correlations violating (EM4) – the concern raised by Liberati and Maccione; (c) whether future ensemble GRB data show $\sqrt{D}$-scaling spread, D-scaling, or a null result. Each channel yields a definite, testable outcome.

6.5 Addendum to Theorem 1: an explicit finite-N bound

Theorem 1 states that decorrelation forces $\sigma\Phi \sim \sqrt{N}$ rather than $\Phi N \propto N$, but stops at the scaling relation. Given only pairwise uncorrelatedness and a common finite variance $\mathrm{Var}(\delta\varphi) = v$ (no independence or specific distribution assumed - exactly the hypotheses already in (EM4)), Chebyshev's inequality applied to $\mathrm{Var}(\Phi N) = Nv$ gives, for every $k > 0$,

$$P(\,|\Phi^N| > k\sqrt{Nv}\,) \leq 1/k^2,$$

a concrete, finite-N, distribution-free probabilistic guarantee (not merely an asymptotic or qualitative statement) that the accumulated phase stays within $O(\sqrt{N})$ of zero with quantified confidence at any fixed N, e.g. taking k=10 bounds the probability of a 10σ or larger coherent-looking excursion at 1%, for any N, under (EM4) alone.

7. Observational Signature: What Astronomers Can Look For

The theorem establishes that no coherent, linear-in-E dispersion signal is expected from the class of emergent-metric discrete theories satisfying the decorrelation condition. However, the residual statistical phase variance $\sigma_\Phi \sim \sqrt{N}$ is not zero: it represents a genuine, in-principle detectable signature. This section specifies what form that signature takes, distinguishing what can be said model-independently from what requires a specific theory.

7.1 The Universal Model-Independent Formula

For any theory in the emergent-metric class satisfying the conditions of the theorem, the two cases are:

$$\Delta t_{coh} \sim \left(\frac{E}{E_{Pl}}\right) \cdot \left(\frac{D}{c}\right) \qquad \textit{[coherent, fixed-lattice case]}$$

$$\sigma(\Delta t) \sim \left(\frac{E}{E_{Pl}}\right) \cdot \frac{\sqrt{\ell_{eff} \cdot D}}{c} \qquad \textit{[decorrelated, emergent-metric case]}$$

where D is the source distance, $\frac{E}{E_{Pl}} \sim 10^{-17}$ for high-energy GRB photons, and ℓ_{eff} is the effective step scale of the discrete substrate is a free parameter at this level of generality, to be fixed by the specific theory. The decorrelated signal grows as $\sqrt{D}$ rather than as D (the coherent case) or as a constant (no effect). The theorem therefore implies a universal $\sqrt{D}$ scaling for stochastic arrival-time spreading, once the number of statistically independent reconstruction domains is proportional to propagation distance.

7.2 What Requires a Specific Theory

Three inputs are needed to convert σ(Δt) into a definite number:

(i) The value of $\boldsymbol{\ell_{eff}}$. In any specific model this is fixed by the theory's own calibration procedure - for instance, by requiring the spectral dimension $d_s \to 4$ at large scales, or by matching to the observed cosmological constant. Without this, ℓ_{eff} is a free parameter ranging from the Planck length upward.

(ii) The correlation length ξ of residual phase correlations. The decorrelation condition $\langle \delta\varphi_i \delta\varphi_j \rangle \approx 0$ is an idealization. Realistic theories will have a finite correlation length ξ such that $\langle \delta\varphi_i \delta\varphi_j \rangle \sim \exp\left(-\frac{|i-j|}{\xi}\right)$. When $\xi > 1$ (in units of ℓ_{eff}), the effective step size in the $\sqrt{N}$ formula becomes $\ell_{eff} \cdot \xi$ rather than ℓ_{eff} alone. The value of ξ is a dynamical property of the specific reconstruction procedure R.

(iii) The angular dependence. The residual signal is statistical and manifests as a correlation between arrival-time spreads of GRBs in nearby sky directions. The angular profile of this correlation depends on whether the theory has large-scale anisotropy in its ensemble structure. A fully isotropic emergent-metric theory predicts no preferred direction in the sky correlation; a theory with residual anisotropy predicts a direction-dependent enhancement. This is a model-specific, not a universal, prediction.

7.3 Observational Strategy

The following observational programme follows directly from the theorem, independent of any specific model:

Do not look for a systematic delay in a single GRB. The theorem rules out a coherent, linear-in-E delay at the level of current sensitivity. A detection in a single source would falsify the decorrelation condition, not confirm the emergent-metric framework.

Look for $\sqrt{D}$-scaling of the arrival-time spread across a large ensemble of GRBs. Collect $\sigma(\Delta t)$ for many sources at different distances D and fit to $\sigma(\Delta t) \propto \sqrt{D}$. A detection of this scaling, distinct from both $\sigma \propto D$ (coherent LIV) and σ = const (no effect), would be positive evidence for the emergent-metric class as a whole.

Look for sky-direction correlations in the spread, not in the mean. Sources in nearby sky directions should show correlated σ(Δt) values if the residual phase correlations have a spatial structure in the emergent geometry. Sources in orthogonal directions should be uncorrelated. This is the angular analogue of the $\sqrt{N}$ formula.

These signatures are within the projected reach of next-generation observatories (CTA, AMEGO-X, future space-based GRB monitors) for sufficiently large source catalogues, even though no single burst will show a detectable effect. The key quantity is not the per-burst sensitivity but the *ensemble* statistical power across hundreds or thousands of GRBs at known distances- a regime that becomes accessible as GRB catalogues grow.

Summary. The theorem predicts a specific, falsifiable observational signature: $\sqrt{D}$-scaling of the stochastic arrival-time spread in GRB ensembles, with no systematic per-energy delay in individual sources. The amplitude of this signal requires a specific theory to fix ℓ_{eff} and ξ; the scaling law itself is model-independent and follows from the decorrelation condition alone. A detection of coherent linear delay would falsify the decorrelation condition. A detection of $\sqrt{D}$-scaling spread would support the emergent-metric class. An exact null result in both channels would constrain $\ell_{eff} \cdot \xi$ to be exponentially small, placing a strong bound on any specific model in this class.

Appendix I

Classification of Discrete Spacetime Theories Under the Emergent-Metric Theorems

This appendix classifies known discrete spacetime theories according to whether they satisfy conditions (EM1)–(EM4) of the emergent-metric class, and are therefore covered by the Main Theorem and its two component theorems. The classification is organised into three groups: theories that satisfy (EM1)–(EM4) and fall within the theorem's scope; theories that violate at least one condition and fall outside it; and borderline cases requiring model-specific analysis. For each theory we indicate which conditions are satisfied, which are open questions, and which are violated, together with the primary references.

A. Theories Satisfying (EM1)–(EM4): Within Scope of the Main Theorem

For theories in this group, the Main Theorem applies: linear Lorentz-violating dispersion is structurally absent in all inertial frames, and the residual stochastic signal scales as $\sqrt{D}$ across a GRB ensemble.

A1. Causal Set Theory (Bombelli, Lee, Meyer, Sorkin 1987; Sorkin 2005). The discrete substrate is a partially ordered set of events, with the metric entirely absent at the fundamental level and reconstructed statistically from the causal order. (EM1): satisfied by definition - no primitive metric. (EM2): satisfied - the Lorentz-invariant Poisson sprinkling has no canonical embedding. (EM3): satisfied - all observables are defined via the causal order. (EM4): proven explicitly by Dowker, Henson & Sorkin (2004); statistical independence of sprinkled points implies decorrelation. To our knowledge, Causal Set Theory is currently the only discrete spacetime framework in which a property equivalent to EM4 has been explicitly demonstrated. Boost invariance of this decorrelation property is additionally established directly (Theorem 2, §6.1, Lemma D), since the Poisson sprinkling is Poincaré-invariant in law and standard causal-order-built reconstruction statistics are causal-order-equivariant.

A2. Causal Dynamical Triangulations (CDT) (Ambjørn, Jurkiewicz, Loll 1998–present). Spacetime is constructed as a Monte Carlo ensemble of simplicial triangulations with a causal constraint. The physical metric is reconstructed from the ensemble average; individual triangulations have no preferred orientation. (EM1): satisfied - metric is a derived, ensemble-averaged quantity. (EM2): satisfied - no background embedding in the ensemble. (EM3): satisfied - observables are ensemble averages. (EM4): open question. The decorrelation of local phase contributions along reconstructed geodesics has not been computed explicitly. The spectral dimension result d_s → 4 in the infrared is consistent with (EM4) but does not prove it. Status: within scope conditionally on (EM4); explicit verification needed.

A3. Loop Quantum Gravity / Spinfoam Models (Rovelli, Reisenberger, Perez, Engle et al.). Geometry is encoded in spin-network states and evolves through spinfoam amplitudes without a background metric. (EM1): satisfied - metric operators are derived. (EM2): satisfied in the canonical and covariant formulations. (EM3): satisfied - observables are defined relative to relational, diffeomorphism-invariant quantities. (EM4): open question, and subject to the concern raised by Liberati & Maccione (2011, 2014) that radiative corrections in the low-energy effective field theory may generate Lorentz-violating operators that produce long-range phase correlations. Status: within scope in principle; (EM4) is the critical open question for LQG phenomenology.

A4. Asymptotic Safety Gravity (Reuter 1998; Reuter & Saueressig 2012). The gravitational path integral is defined via a non-perturbative renormalisation group flow, with the metric as a dynamical integration variable over all geometries. No fixed background is singled out in the background-field formalism (the background is a gauge choice, not a physical preferred frame). (EM1)–(EM3): satisfied in the background-independent formulation. (EM4): open question; depends on whether the specific fixed-point effective action induces anisotropic correlations in the propagating sector. Status: within scope conditionally on (EM4).

A5. Group Field Theory (Oriti 2006–present; Gielen, Sindoni). Spacetime is not assumed; it emerges from the condensation of combinatorial 'atoms of space' described by a group field. (EM1)–(EM3): satisfied by construction - there is no background spacetime at the fundamental level. (EM4): not investigated in the literature. Status: within scope in principle; (EM4) is open and constitutes a concrete research question for GFT phenomenology.

A6. Relative Locality (Amelino-Camelia, Freidel, Kowalski-Glikman, Smolin 2011). Spacetime locality is observer-dependent; the fundamental arena is a curved momentum space, and spacetime emerges as a derived concept. (EM1)–(EM3): plausibly satisfied, as there is no fixed background coordinate spacetime. (EM4): not investigated. Note: relative locality is primarily a kinematic framework, not a full dynamical theory, so the status of (EM4) as a property of a 'reconstruction procedure R' is not straightforwardly defined. Status: partially within scope; requires translation into the (G, R) language of the theorem.

B. Theories Violating (EM1)–(EM4): Outside Scope of the Main Theorem

For theories in this group, the Main Theorem does not apply. Linear Lorentz-violating dispersion may or may not be present, depending on the specific model, but the structural protection established by the theorem is absent. GRB and UHECR constraints from the Jacobson–Liberati–Mattingly programme apply directly to this group.

B1. Lattice QFT / Hypercubic Lattice Models (Wilson 1974 and standard lattice formulations). A fixed hypercubic lattice with constant spacing ℓ provides a preferred orientation at the Planck scale. (EM2) is violated: the embedding is canonical and fixed. Rotational and Lorentz symmetry are recovered only approximately in the continuum limit, with corrections of order $\left(\frac{\ell}{L}\right)^n$. Generic expectation: linear LIV is present at order $\frac{\ell}{L_{macro}}$, consistent with the Corollary of the Main Theorem. Note: this is the dominant implicit assumption in much of the GRB phenomenology literature when 'discrete spacetime' is discussed.

B2. Hořava Gravity (Hořava 2009). Explicit anisotropy between space and time is built into the action via a preferred foliation structure. (EM2) is violated by the fixed foliation, and (EM3) is violated since local observables are defined relative to the preferred time direction. Linear Lorentz violation is a generic prediction in the UV, with suppressed corrections in the IR.

B3. Standard Model Extension (SME) (Colladay & Kostelecký 1997, 1998). By construction, Lorentz-violating operators are added to the Standard Model action with fixed, observer-independent background tensor fields (e.g. $k_{\mu\nu\rho\sigma}$). (EM2) is violated: these tensors define preferred directions in spacetime. The SME is the primary phenomenological framework for LIV constraints and is entirely outside the emergent-metric class. The GRB bounds of Jacobson, Liberati & Mattingly apply directly to SME coefficients.

B4. Rainbow Gravity (Magueijo & Smolin 2004). The metric depends on the energy of the probe particle, with a fixed preferred frame (usually the CMB rest frame). (EM2) is violated by the fixed preferred frame. Dispersion is by construction energy-dependent; the theory is designed to produce exactly the type of signal that the Main Theorem shows is absent for the emergent-metric class.

B5. Non-Commutative Geometry with Fixed $\boldsymbol{\theta^{\mu\nu}}$ (Seiberg & Witten 1999 and standard NCQFT). When the non-commutativity parameter $\theta^{\mu\nu}$ is a fixed constant tensor (the most common formulation), it defines a preferred set of directions in spacetime, violating (EM2). This produces Lorentz-violating dispersion at order θE^2. Status: outside scope when $\theta^{\mu\nu}$ is fixed.

C. Borderline Cases Requiring Model-Specific Analysis

These theories require explicit investigation to determine whether (EM1)–(EM4) are satisfied. The Main Theorem may apply, but this must be verified case by case.

C1. String Theory / M-Theory (in the background-independent formulation). M-theory is believed to be background-independent in principle, but almost all practical computations are performed on fixed string backgrounds (flat space, anti-de Sitter, etc.), which violate (EM2). In the genuine background-independent formulation (if and when it is established), (EM1)–(EM3) would plausibly be satisfied. (EM4) has not been investigated. Status: outside scope in perturbative formulations; potentially within scope in a background-independent formulation, but that formulation does not currently exist in explicit form.

C2. Doubly Special Relativity (DSR) (Amelino-Camelia 2001, 2002; Magueijo & Smolin 2002). DSR modifies Lorentz transformations to make the Planck energy an observer-independent invariant, rather than constructing a background-independent substrate. It is a kinematic framework, not a theory of a discrete substrate. (EM1)–(EM3) are not straightforwardly applicable – DSR now does not specify a reconstruction procedure R. DSR does not predict linear dispersion in the standard sense (it deforms the dispersion relation but preserves a modified relativity principle). Status: not directly covered by the theorem; the theorem is stated for theories with an underlying substrate G and reconstruction R, which DSR does not provide. DSR is neither a counterexample nor a confirmation.

C3. Non-Commutative Geometry with Dynamical $\boldsymbol{\theta^{\mu\nu}}$ (some formulations). If the non-commutativity tensor is a dynamical field rather than a fixed background, (EM2) may be satisfied. (EM4) depends on the dynamics of $\theta^{\mu\nu}$. Status: open; requires analysis of the specific dynamical model.

C4. TQGT* (Materov 2025–present). A Z^4 hypercubic torus with metric reconstructed via spectral dimension ($d_s \to 4$ for $L \geq 5$) and plaquette Hamiltonians supporting U(1)/SU(2)/SU(3) gauge groups. (EM1): satisfied - no primitive metric. (EM2): satisfied in the dynamically reconfiguring graph construction; the fixed Z^4 topology provides a combinatorial structure, not a fixed metric embedding. (EM3): satisfied - observables are defined via spectral and diffusion quantities. (EM4): open question. The decorrelation of local phase contributions along reconstructed geodesics has not been computed explicitly; this is identified as a target computation for Quantumograph v14. Status: within scope in principle; (EM4) is the primary open verification task for this theory.

D. Summary Table

The following table summarises the classification.
✓ = condition satisfied,
? = open question,
✗ = condition violated.

Theory	EM1	EM2	EM3	EM4	Status
Causal Sets	✓	✓	✓	✓	Within scope (proven)
CDT	✓	✓	✓	?	Within scope (EM4 open)
LQG / Spinfoams	✓	✓	✓	?	Within scope (EM4 open)
Asymptotic Safety	✓	✓	✓	?	Within scope (EM4 open)
Group Field Theory	✓	✓	✓	?	Within scope (EM4 open)
Relative Locality	✓	✓	?	?	Partially within scope
TQGT*	✓	✓	✓	?	Within scope (EM4 open)
Lattice QFT	✓	✗	✗	✗	Outside scope

Hořava Gravity	✓	✗	✗	✗	Outside scope
SME	✓	✗	✗	✗	Outside scope
Rainbow Gravity	✓	✗	✗	✗	Outside scope
NCQFT (fixed θ)	✓	✗	✗	✗	Outside scope
String Theory (pert.)	✓	✗	✗	✗	Outside scope
DSR	n/a	n/a	n/a	n/a	Not applicable

*Quantumograph included only as an illustrative example of a theory for which EM4 remains to be established. This illustrates that the present theorem is not tailored to any specific model, including the author's own framework.

AI Disclosure: My custom AI assistant was used solely as a processing tool for preliminary literature search and data clustering in compiling this classification table, as well as "editor" for editing the text into advanced academic English. All criteria selection, physical data verification, and theoretical conclusions were performed exclusively by the author.

References

Carlip, S. Spacetime Foam: A Review. arXiv:2209.14282 [gr-qc] (2022)

Amelino-Camelia, G. “Quantum-Spacetime Phenomenology.” Living Reviews in Relativity 16 (2013): 5.

Amelino-Camelia, G. “Doubly-Special Relativity: First Results and Key Open Problems.” International Journal of Modern Physics D 11 (2002): 1643. arXiv:gr-qc/0210063.

Amelino-Camelia, G., Freidel, L., Kowalski-Glikman, J., and Smolin, L. “The Principle of Relative Locality.” Physical Review D 84 (2011): 084010. arXiv:1101.0931.

Abdo, A. A., et al. (Fermi-LAT and Fermi-GBM Collaborations). “A Limit on the Variation of the Speed of Light Arising from Quantum Gravity Effects.” Nature 462 (2009): 331–334.

Vasileiou, V., Jacholkowska, A., Piron, F., et al. “Constraints on Lorentz Invariance Violation from Fermi-Large Area Telescope Observations of Gamma-Ray Bursts.” Physical Review D 87 (2013): 122001.

MAGIC Collaboration (Acciari, V. A., et al.). “Bounds on Lorentz Invariance Violation from MAGIC Observation of GRB 190114C.” Physical Review Letters 125 (2020): 021301.

H.E.S.S. Collaboration (Abdalla, H., et al.). “The 2014 TeV γ-Ray Flare of Mrk 501 Seen with H.E.S.S.: Temporal and Spectral Constraints on Lorentz Invariance Violation.” The Astrophysical Journal 870 (2019): 93.

Vasileiou, V., et al. “A Planck-Scale Limit on Spacetime Fuzziness and Stochastic Lorentz Invariance Violation.” Nature Physics 11 (2015): 344–346.

Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. D. “Space-Time as a Causal Set.” Physical Review Letters 59 (1987): 521–524.

Sorkin, R. D. “Causal Sets: Discrete Gravity.” Lectures on Quantum Gravity, edited by A. Gomberoff and D. Marolf, Springer (2005): 305–327. arXiv:gr-qc/0309009.

Dowker, F., Henson, J., and Sorkin, R. D. “Quantum Gravity Phenomenology, Lorentz Invariance and Discreteness.” Modern Physics Letters A 19 (2004): 1829–1840. arXiv:gr-qc/0311055.

Ambjørn, J., Jurkiewicz, J., and Loll, R. “The Universe from Scratch.” Contemporary Physics 47 (2006): 103–117. arXiv:hep-th/0509010.

Rovelli, C. Quantum Gravity. Cambridge University Press, 2004.

Perez, A. “The Spin Foam Approach to Quantum Gravity.” Living Reviews in Relativity 16 (2013): 3.

Reuter, M. “Nonperturbative Evolution Equation for Quantum Gravity.” Physical Review D 57 (1998): 971–985. arXiv:hep-th/9802150.

Reuter, M. and Saueressig, F. Quantum Einstein Gravity. New Journal of Physics 14 (2012): 055022.

Oriti, D. “The Group Field Theory Approach to Quantum Gravity.” In Approaches to Quantum Gravity, edited by D. Oriti, Cambridge University Press (2009). arXiv:gr-qc/0607032.

Colladay, D. and Kostelecký, V. A. “Lorentz-Violating Extension of the Standard Model.” Physical Review D 58 (1998): 116002. arXiv:hep-ph/9809521.

Hořava, P. “Quantum Gravity at a Lifshitz Point.” Physical Review D 79 (2009): 084008. arXiv:0901.3775.

Magueijo, J. and Smolin, L. “Gravity’s Rainbow.” Classical and Quantum Gravity 21 (2004): 1725–1736. arXiv:gr-qc/0305055.

Jacobson, T., Liberati, S., and Mattingly, D. “Threshold Effects and Planck Scale Lorentz Violation: Combined Constraints from High-Energy Astrophysics.” Physical Review D 67 (2003): 124011. arXiv:hep-ph/0209264.

Jacobson, T., Liberati, S., and Mattingly, D. “Lorentz Violation at High Energy: Concepts, Phenomena and Astrophysical Constraints.” Annals of Physics 321 (2006): 150–196. arXiv:astro-ph/0505267.

Liberati, S. and Maccione, L. “Lorentz Invariance Violation: Fields and Gravity.” Physical Review Letters 112 (2014): 151301. arXiv:1309.7296.

Seiberg, N. and Witten, E. “String Theory and Noncommutative Geometry.” Journal of High Energy Physics 9 (1999): 032. arXiv:hep-th/9908142.

Basu, S. and Mattingly, D. “Constraints from Cosmic Rays on Non-Systematic Lorentz Violation.” Classical and Quantum Gravity 22 (2005): 3029–3038. arXiv:astro-ph/0501425.

Mattingly, D. “Modern Tests of Lorentz Invariance.” Living Reviews in Relativity 8 (2005): 5.

Materov, S. “Quantumograph: A Testable Quantum Graph Theory of Spacetime v13”, (2026). https://doi.org/10.5281/zenodo.18410313

Materov, S. Computational Appendix - Boost Invariance of Decorrelation (Section 6). GitHub repository, 2026. https://github.com/SergejMaterov/NoDispersion.
