The Quantum-Geometric Correspondence: Status and Outlook

A review of the axioms, the gravitational-decoherence programme, and its falsifiable predictions

Marc Sperzel*

Independent Researcher
MSci Physics, King's College London

July 2, 2026

Abstract

The Quantum-Geometric Correspondence treats quantum mechanics and general relativity as complementary projections of a single quantum-informational reality. We review the programme's status. Three primitive axioms—generalized entropy, entanglement equilibrium, and modular time—recover the Bekenstein–Hawking entropy, the holographic bound, the Unruh temperature, and the Einstein equations, and yield a thermodynamic dark-energy density $\rho_{\mathrm{DE}} = \alpha c^2 H^2/G$. The programme's one new dynamical prediction is gravitational decoherence of a mass superposition at the rate $\Gamma_{\mathrm{dec}} = C\,E_G/\hbar$ with $E_G = GM^2/d$ and $C \in [2/\pi, 1]$, giving $\tau_{\mathrm{dec}} \approx 1.6$ ns for a microgram over a millimetre (the point-mass benchmark; an extended object's self-energy saturates at $1.2\,GM^2/R$)—some thirty-four orders of magnitude faster than the perturbative ($\propto G^2$) result, and hence experimentally decisive. We distinguish what is established from what is not: the energy scale E_G is proven, while the operator identity $K = 2\pi H_{\mathrm{phys}}$ that converts it into a rate is a conjecture reduced, through a finite causal diamond (Type II crossed-product) construction, to a single open question—the clock-model normalization that fixes C. We collect the falsifiability map: a picogram silica superposition held for one second in cryogenic extreme-high vacuum gives a visibility drop to 0.18 versus the standard ≈ 1, decided in roughly thirty interference runs. We summarize the cosmological sector (the MOND scale $a_0 = cH_0/(2\pi)$, holographic dark energy, and a cosmic decoherence rate $\Gamma_{\mathrm{total}} = g(1)H/(4\pi)$ that share a single de Sitter frequency), position the programme against spontaneous-collapse, Diósi–Penrose, and stochastic-gravity approaches, and give the open problems and roadmap. The balance is candid: the physics and the falsifiable claim are intact, the central identity carries the status of a strong conjecture with no experimental confirmation, and an experiment this decade can decide it.

Contents

*ORCID: 0009-0000-6252-3155. Email: me@marcsperzel.com

1 Introduction and scope

Quantum mechanics and general relativity are mutually incompatible in the regime where both matter, yet each is among the best-tested theories in physics. The Quantum-Geometric Correspondence (QGC) takes a conservative stance on this incompatibility: it neither quantises the metric nor postulates a discrete substrate, but treats the two theories as the matter-superposition reading and the geometry-entanglement reading of one underlying quantum state [1], related by a change of observational framework rather than by a dynamical approximation.

This review states what the programme has established, what it has not, and what would decide it. Its central energy scale is on firm ground; its headline mechanism is a well-posed conjecture reduced to a single open construction; its sharpest prediction is falsifiable but not yet tested. Throughout, claims are labelled as established (rigorous proof or exact numerical identity), derived (following from the stated axioms under explicitly named assumptions), physically motivated (argued and consistent with bounds, but not derived), or conjectured (supported in special cases, not established in general). No claim in this review carries experimental confirmation.

1.1 The two faces of the correspondence

Three primitive axioms (Section 2)—generalized entropy, entanglement equilibrium (whose stationarity returns the Einstein equations and thermal equilibrium together), and modular time (which identifies the diamond's modular flow with physical time)—recover the Bekenstein–Hawking entropy, the holographic bound, and the Unruh temperature, and yield one new structural claim: placing matter in a spatial superposition and entangling spacetime geometry are dual descriptions of one process. The direct physical consequence is *gravitational decoherence*—a spatial superposition of mass loses coherence because the dual geometric description is an entangling channel.

1.2 The central quantitative claim, and its honest status

For a mass M in a superposition of two locations separated by d, the programme's decoherence time is

$$\tau_{\mathrm{dec}} \;=\; \frac{\hbar d}{GM^2} \;=\; \frac{\hbar}{E_{\mathrm{G}}}, \qquad E_{\mathrm{G}} = \frac{GM^2}{d}, \tag{1}$$

the Diósi–Penrose form [9, 10], with $\tau_{\mathrm{dec}}(1\,\mu\mathrm{g}, 1\,\mathrm{mm}) \approx 1.6$ ns. The *energy scale* $E_{\mathrm{G}} = GM^2/d$ is established: it is the classical gravitational self-energy difference between the branches, and nothing in the programme is needed to obtain it. We conjecture that the decoherence *rate* equals this energy scale divided by $\hbar$,

$$\Gamma_{\mathrm{dec}} \;=\; C\,\frac{E_{\mathrm{G}}}{\hbar}, \qquad C \in [\,2/\pi,\,1\,] \approx [0.637, 1.000], \tag{2}$$

the coefficient C bounded but not uniquely pinned (Section 3). This is the programme's load-bearing prediction: perturbative quantum field theory gives a rate suppressed by an additional power of G—a decoherence time of order 10^{18} years for the same parameters, some 34 orders of magnitude longer. The two predictions are operationally distinct—one says a microgram superposition decoheres in nanoseconds, the other says it never decoheres in any laboratory—and an experiment in the right regime (Section 4) decides between them.

1.3 The structural development underlying the rate

The operator identity underlying (2)—that the region's modular Hamiltonian equals 2π times the constrained physical Hamiltonian, $K = 2\pi H_{\text{phys}}$—cannot be established in flat space, where the relevant local algebra is of Type III_1 and admits no trace. Relocating the construction to a finite causal diamond gives an algebra of Type II, to which the von Neumann machinery of Chandrasekaran–Longo–Penington–Witten [21] applies. A sequence of results (Section 3) assembles that finite-diamond construction into a single object whose one remaining free parameter is the coefficient C of (2). Whether that construction can be closed at full rigour is the question on which the operator identity—and with it the standing of the programme's central claim—turns.

1.4 Plan of the review

Section 2 states the three axioms and the theorems they generate. Section 3 develops the gravitational-decoherence core: the mechanism, the rate, the G^1-versus-G^2 distinction, and the finite-diamond construction. Section 4 collects the testable predictions and the experimental falsifiability map. Section 5 turns to the cosmological sector—the MOND acceleration scale, holographic dark energy, and cosmic decoherence—which shares the programme's de Sitter thermodynamics. Section 6 positions QGC against spontaneous-collapse models, the Diósi–Penrose proposal, and stochastic gravity. Section 7 states the open problems and a roadmap, and Section 8 concludes. The underlying computations are carried in the cited papers.

2 Axiomatic foundation

The canonical formulation (v3.0) rests on three primitive axioms—the algebra, state, and time clauses of one statement: the physical content of a region is the observer-dressed modular structure of its causal diamond. The broader structure that earlier formulations carried as six axioms plus a duality correspondence, and the Einstein equations themselves, are recovered here as derived theorems.

2.1 The three primitive axioms

Definition 2.1 (Generalized entropy, Axiom I)**.** For a region bounded by a quantum extremal surface or causal horizon $\mathcal{X}$, with the gravitational constraints imposed, the physical observables form a Type II (crossed-product) algebra whose von Neumann entropy is the generalized entropy

$$S_{\text{gen}}(\mathcal{X}) \;=\; \frac{A(\mathcal{X})}{4\ell_{\text{P}}^2} + S_{\text{ext}}(\mathcal{X}) + \text{const}, \tag{3}$$

where A is the proper area of $\mathcal{X}$, $\ell_{\text{P}} = \sqrt{G\hbar/c^3}$ is the Planck length, and S_{ext} is the von Neumann entropy of the exterior fields. For the total state on a complete Cauchy slice S_{gen} is invariant under evolution (global unitarity); for open subregions the generalized second law $\Delta S_{\text{gen}} \geq 0$ holds along future horizon cuts. Only the sum, its differences, and its monotonicity are physical; the area law is the framework's irreducible geometric input, postulated here, not derived.

Definition 2.2 (Entanglement equilibrium, Axiom II)**.** Every small causal diamond D ($\ell_\mathrm{P} \ll \ell \ll L_\mathrm{curv}$) is at entanglement equilibrium: the generalized entropy of Axiom I is stationary at fixed volume,

$$\delta S_\mathrm{gen}(D) \;=\; 0. \tag{4}$$

Stationarity, together with the first law of entanglement and the Raychaudhuri equation, yields the semiclassical Einstein equations sourced by $\langle T_{\mu\nu}\rangle$—Jacobson's thermodynamic derivation of the field equations [27], run on the generalized entropy with no Einstein–Hilbert term inserted by hand. The equilibrium state satisfies the KMS condition at $\beta_\mathrm{mod} = 2\pi$. The entropic free-energy functional $S[\rho, g] = \int d^4x\, \sqrt{-g}\,[\mathrm{Tr}(\rho\hat{H}) + c^4 R/16\pi G - s(\rho)/\beta]$ that earlier formulations carried as this axiom survives as an equivalent repackaging of the same equilibrium statics, not as a primitive postulate.

Definition 2.3 (Modular time, Axiom III)**.** For the diamond's observer, the modular flow of the constrained (Type II, observer-dressed) algebra at the KMS temperature $\beta_\mathrm{mod} = 2\pi$ *is* physical time evolution; at the operator level,

$$K \;=\; 2\pi H_\mathrm{phys} + \mathcal{O}(G^2). \tag{5}$$

This fixes the G^1 decoherence rate $\Gamma_\mathrm{dec} = C\, E_\mathrm{G}/\hbar$ with $E_\mathrm{G} = GM^2/d$ and $C \in [2/\pi, 1]$ (natural value $C = 1$, the Diósi–Penrose rate). The identity is proven at the level of expectation values and in solvable models and conjectured as a full operator identity; it is the framework's falsifiable clause. Axiom III supersedes the earlier "Gravitational Information Saturation Principle": the Margolus–Levitin reading of that principle supplies the $C = 2/\pi$ floor of the same rate.

2.2 Derived theorems

Axioms I–III recover the following standard results, each established or derived as marked, with dimensional consistency and limiting behaviour checked.

- **Bekenstein–Hawking entropy** $S_\mathrm{BH} = A/(4\ell_\mathrm{P}^2)$ [28] and the **holographic bound**: established from Axiom I.
- **Einstein equations** $G_{\mu\nu} = (8\pi G/c^4)\langle T_{\mu\nu}\rangle$: derived from Axioms I and II jointly (Jacobson's theorem).
- **Unruh temperature** $T_U = \hbar a/(2\pi c k_\mathrm{B})$ [29]: established from Axiom I via the horizon thermodynamics.
- **Duality correspondence**: a matter superposition and the corresponding geometric entanglement are dual descriptions, derived from Axioms I and II. This is the structural claim that drives gravitational decoherence (Section 3).
- **Dark energy** $\rho_\mathrm{DE} = \alpha\, c^2 H^2/G$ with $\alpha \approx 0.082$, derived from Axioms I and II [6] (Section 5).
- **Observer-dependent horizons**: physically motivated from Axiom I—a dimensional-analysis relation between acceleration and horizon scale, not a derivation (the boxed coefficient absorbs an $\mathcal{O}(1)$ factor).

2.3 Demoted: the Planck-scale interpolation

Earlier formulations carried a fourth statement—a Planck-scale interpolation $\mathcal{O}_{\text{unified}} = \rho_{\text{quantum}}\, f(r/\ell_{\text{P}}) + \rho_{\text{geometric}}[1-f]$, $f(x) = 1/(1+x^2)$—as a primitive axiom. It is not one here: wherever the Einstein equations hold the two densities coincide and the interpolation is trivial, and where they are asserted to fail the interpolated source violates the contracted Bianchi identity and the radial profile singles out a frame. What survives is a speculative ultraviolet extension—the minimum length $\Delta x_{\min} \sim \sqrt{2}\,\ell_{\text{P}}$ and the modified dispersion—on the generalized-uncertainty route, with explicit $\mathcal{O}(1)$-coefficient and power-law uncertainty rather than as an axiom consequence.

2.4 What the axioms do and do not fix

Axioms I and II fix the kinematics (what the dual descriptions are), the equilibrium dynamics (the Einstein equations), and the local state (the Gibbs form). Axiom III supplies the one genuinely new dynamical prediction, the G^1 rate, through the operator-level identity between a region's modular Hamiltonian and its constrained physical Hamiltonian. In four-dimensional linearized gravity the several heuristic arguments once advanced for G^1 scaling reduce to a single physical input—the Wheeler–DeWitt Hamiltonian constraint—rather than to several independent derivations. The axiomatic layer is internally consistent and recovers known physics; the operator identity of Axiom III remains conjectural, and it is exactly what experiment tests.

3 Gravitational decoherence: mechanism, rate, and the finite-diamond construction

Gravitational decoherence is the core of QGC and the source of its sharpest prediction. This section separates what is established from what is conjectured.

3.1 The energy scale

For a mass M in a superposition of two configurations differing by a separation d, the gravitational self-energy difference between the branches is

$$E_{\text{G}} = \frac{GM^2}{d} = 6.7\times10^{-26}\ \text{J} \quad (1\,\mu\text{g},\ 1\,\text{mm}). \tag{6}$$

For an extended source the point-mass form is replaced by the exact Diósi–Penrose self-energy integral, which saturates at $1.2\,GM^2/R$ for separations beyond the object diameter. The energy scale (6) is established: it is a statement of classical gravity, independent of any QGC-specific machinery, and it sets the only scale that can appear in a gravitationally induced rate.

3.2 The two scalings, G^1 versus G^2

The physical question is whether the decoherence rate is $\Gamma_{\text{dec}} \sim E_{\text{G}}/\hbar \propto G^1$ or the perturbative-QFT result $\propto G^2$. The distinction has a clean interpretation, confirmed by explicit computation:

- the G^2 rate is the rate of *entanglement generation* between matter and the graviton field, computed perturbatively from an unphysical (un-constrained) initial state;

- the G^1 rate is the rate of *entanglement manifestation* in the constrained theory, where the Wheeler–DeWitt Hamiltonian constraint ties the gravitational field to its source.

The G^2 rate correctly describes generation from the unphysical product state, but because that state violates the constraint, the framework's physical prediction is the G^1 rate; the two are a falsifiable fork, not co-equal answers. For $(1\,\mu\mathrm{g}, 1\,\mathrm{mm})$ they differ by a factor $\sim 3 \times 10^{34}$: $\tau_{\rm dec} \approx 1.6$ ns (G^1) versus $\tau_{G^2} \sim 10^{18}$ yr. The G^1 *scaling* is derived in linearized gravity by a coherent-state argument; the full spin-2 computation confirms the scalar-proxy result [2], and soft- graviton infrared logarithms do not promote G^2 to G^1. What remains conjectured is the operator-level identity that converts the energy scale into the rate.

3.3 The operator identity and the flat-space obstruction

In the modular-theory formulation [1], the decoherence rate follows from the identity

$$K \;=\; 2\pi\, H_{\rm phys} + \mathcal{O}(G^2), \tag{7}$$

where K is the modular Hamiltonian of the region containing the superposition and $H_{\rm phys}$ is the constrained physical Hamiltonian.

Conjecture 3.1 (Modular–physical identity)**.** *The identity* (7) *holds as an operator equation. It is established in expectation values and in solvable models.*

Two obstructions prevent a flat-space proof. First, the local algebra of a region in the vacuum is a Type III_1 von Neumann factor: it has no trace, no density matrices, and no canonical "rate per unit time" normalization. Second, the Bisognano–Wichmann [18] transport of the vacuum modular Hamiltonian to the matter-loaded, gravitationally dressed state fails because the Newtonian field has global support, lying outside the region's algebra. Identity (7) therefore remains conjectural in flat space.

3.4 The finite-diamond construction

The route around the obstruction is to give the region a cosmological constant—equivalently, to work in a finite causal diamond—where the algebra becomes Type II and carries a trace, and the crossed-product machinery of Chandrasekaran–Longo–Penington–Witten [21] and Jensen–Sorce–Speranza [22] applies. The construction assembles six results into a single object.

1. **Type II trace, unconditionally.** By Takesaki's structure theorem [19], the crossed product of the diamond algebra by its modular automorphism group is Type II_∞ with a canonical semifinite trace, for *any* modular flow. Established.

2. **Geometric flow for conformal matter.** The Casini–Huerta–Myers conformal map [20] sends the diamond to $\mathbb{R} \times \mathbb{H}^3$, turning the conformal-Killing modular flow into an exact, constant-norm Killing flow. Derived for the conformal proxy.

3. **Geometric flow for the graviton.** The physical graviton is not a conformal field, but its gauge-invariant ball algebra is two scalar towers with $\ell \geq 2$ (Benedetti–Casini [23]), each inheriting the local weight $f(r) = (R^2 - r^2)/(2R)$ (Huerta–van der Velde [24]); the flat-space degeneration of the Regge–Wheeler [25] and Zerilli [26] master potentials to the

scalar centrifugal potential anchors the reduction. The flow is geometric for the physical field. Derived from published inputs.

4. **Intrinsic observer and clock.** The crossed-product observer is the diamond-center geodesic; its modular time is $t(\sigma) = R\tanh(\pi\sigma/2)$, so the diamond tips form a horizon reached only at infinite modular time, while the observer carries a uniform clock at the canonical Kubo–Martin–Schwinger temperature $\beta_{\mathrm{mod}} = 2\pi$. Established exactly.

5. **Residues decoupled from the rate.** Three remaining technical residues—the bounded-below observer Hamiltonian (Type II_1 refinement), the entangling-surface edge modes, and the four-dimensional trace anomaly—are all additive, entropy-constant effects. The modular flow that sets the rate is invariant under $K \to K + c\mathbb{1}$, so none of them can move the coefficient C; they shift only the entropy's additive constant. The shift-invariance is exact.

6. **Infrared safety.** The extracted physical rate is independent of the arbitrary diamond size R: the modular weight at the source ($\propto R$) cancels the observer clock ($\propto R$), giving $\Gamma_{\mathrm{phys}} \propto E_{\mathrm{G}}/\hbar$ with $\partial_R\Gamma_{\mathrm{phys}} = 0$. Exact.

One weight $f(r)$, one Bisognano–Wichmann factor 2π, one modular flow, and one graviton-to-matter degeneration thread run consistently through all six. The assembled construction has a *single* remaining free parameter—the clock-model normalization that fixes the coefficient C—with everything else proven, derived, or scoped.

3.5 One formalism for both the scale and the rate

In earlier formulations the energy scale and the rate were computed in *different* formalisms (linearized gravity for E_{G}, the full constrained theory for the mechanism) and stitched together. The finite-diamond modular Hamiltonian removes the stitch: $K = 2\pi H_{\mathrm{phys}}$ is one object with two observables. Its *gap* between the which-path branches is the energy scale,

$$\frac{\hbar}{2\pi}\,\Delta K \;=\; E_{\mathrm{G}} \qquad \text{(independent of the clock normalization)}, \tag{8}$$

and its *flow* on the observer clock is the rate, $\Gamma_{\mathrm{dec}} = C\,E_{\mathrm{G}}/\hbar$. The energy scale read from the gap is established and independent of the open coefficient; only the rate carries it. The Hamiltonian constraint that the linearized formalism could not make manifest is here the defining identity of the single object.

3.6 The rate coefficient

What the finite-diamond trace fixes is the modular temperature ($\beta_{\mathrm{mod}} = 2\pi$), the variance of the modular Hamiltonian, and the minimum-uncertainty clock width; these remove the normalization freedom that once gave a factor-four window. The residual freedom is a genuine distinction between two operationally different observables built from the same trace-fixed metric: the continuous fringe-visibility (Bures) rate gives $C = 1$ (the Diósi–Penrose value), while the Margolus–Levitin orthogonalization speed gives the floor $C = 2/\pi$. An interferometer measures fringe-visibility decay, so the trace-fixed metric *selects* $C = 1$ for the falsifiable prediction; but the exactness of that selection still rests on a spectral input that the free-field analysis does not supply (the bath is super-Ohmic, with no zero-frequency weight to whiten). The honest statement is therefore

$$C \in [\,2/\pi,\, 1\,], \qquad \text{natural value } C = 1, \text{ not uniquely pinned.} \tag{9}$$

Pinning C to a single value, and proving (7) as an operator equation, are the same open problem—the closure of the finite-diamond construction at full rigour—and it is the problem on which the standing of the programme's central claim turns.

4 Predictions and experimental falsifiability

The programme's value rests on whether it can be tested. The genuinely discriminating prediction is separated here from those that merely accompany it.

4.1 The discriminating prediction

The one prediction that distinguishes QGC from standard physics is the G^1 decoherence rate (2): a spatial superposition of a sufficiently massive object decoheres in (1), whereas perturbative quantum field theory predicts effectively no decoherence ($\tau_{G^2} \sim 10^{18}$ yr). The observable is the interference visibility after a hold time t,

$$V(t) \;=\; e^{-t/\tau_{\mathrm{dec}}}, \qquad \tau_{\mathrm{dec}} = \hbar/E_{\mathrm{G}}, \tag{10}$$

with the standard prediction $V \approx 1$ (Figure 2). A measured visibility decay at the rate (10) would confirm G^1; a null result (sustained coherence well past τ_{dec}) would falsify it [4].

4.2 The falsifiability map

For a uniform sphere of density ρ and radius $R = (3M/4\pi\rho)^{1/3}$ in a superposition of separation Δx, the exact self-energy $E_{\mathrm{G}}(\Delta x)$ saturates at $1.2\,GM^2/R$ for $\Delta x \gtrsim 2R$, giving a decoherence time that scales as $\tau_{\mathrm{dec}} \propto M^{-5/3}$ in the saturated regime. This produces a sharp *threshold*: for silica, $\tau_{\mathrm{dec}} = 1$ s requires

$$M_\star \;\approx\; 7.6 \times 10^{-16}\ \mathrm{kg} \quad (R \approx 0.44\,\mu\mathrm{m}). \tag{11}$$

Below $M_\star$ the superposition decoheres too slowly to test on laboratory timescales; above it, the rate turns on rapidly. The decision is therefore a narrow window in mass and separation, not "the larger the better."

To distinguish $V(t)$ from the standard null $V = 1$ at the 5σ level requires, at the shot-noise floor, $N_{5\sigma} = 25/(1-V)^2$ interference runs. Figure 1 maps the decoherence time over the $(M, \Delta x)$ plane and locates the platforms, and Table 1 reads off the discriminant for each.

The reading is sharp. The picogram envelope target—1.5 pg of silica, a $1\,\mu$m superposition held for 1 s under $\sim 7{\times}10^{-15}$ mbar cryogenic vacuum—is decided in ~ 37 runs. The hold time and superposition scale are at the frontier of demonstrated levitated-optomechanics capability [17], whose ground-state-cooled spheres are some three orders of magnitude lighter; the micrometre superposition is the goal of MAQRO-class proposals [16]. MAQRO's own baseline particle

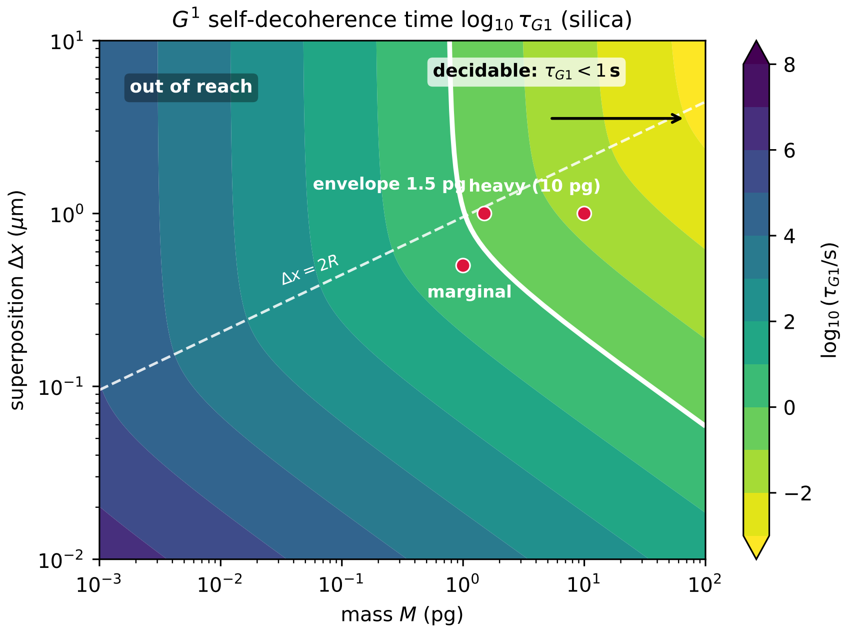

Figure 1. *The G^1 self-decoherence time $\tau_{\rm dec}(M, \Delta x)$ for a silica sphere, from the exact Diósi–Penrose self-energy. The decidable region (decoherence within a one-second hold, $\tau_{\rm dec} < 1$ s) lies to the right of the white contour; the dashed line $\Delta x = 2R$ marks the onset of the saturated regime. Below the threshold mass $M_\star \approx 7.6 \times 10^{-16}$ kg the superposition is out of reach on laboratory timescales. The marked platforms are the rows of Table 1.*

($\sim 10^{-17}$ kg, radius ~ 120 nm), however, lies a factor ~ 50 *below* the threshold $M_\star$: it self-decoheres only on a ~ 600 s timescale and is swamped by residual gas at its design vacuum, so a clean decision requires a $\sim 5\times$ heavier ($\sim$picogram) sphere. Current levitated spheres are sub-threshold and can only set bounds. Large-separation platforms of the Bose–Marletto–Vedral type carry enough mass but, with separations of hundreds of micrometres, sit in the saturated regime and self-decohere in milliseconds—faster than the entanglement protocol they are designed to run, so for them the G^1 signature is a null entanglement result rather than a measured visibility.

4.3 The accompanying predictions, and their honest status

Other consequences are real but less discriminating:

- **Bose–Marletto–Vedral visibility suppression** [14, 15, 3]: an orientation-dependent, approximately mass-independent law, physically motivated. It currently carries an unresolved inconsistency between two candidate self-decoherence rates (the arm-set rate $\propto 1/\Delta$ versus the radius-saturated rate), which differ by $\sim 10^2$ at experimental parameters; the prediction is not firm until that rate is resolved.

- ***N*-particle scaling** $\Gamma_{\rm dec} \propto M_{\rm total}^2$: real, but generic to any mass-density-sourced decoherence, so not unique to QGC.

- **Short-time Zeno regime**: a zero-initial-slope onset $V(t) \approx 1 - \frac{1}{2}M_2 t^2$ for $t < d/c$, distinguishing QGC from a Markovian collapse ansatz but living in a femtosecond window that is not near-term testable.

Platform (representative)	M (kg)	τ_{dec}	$V(t)$	$N_{5\sigma}$
Envelope target (silica, Δx=1 μm, t=1 s)	1.5×10^{-15}	0.59 s	0.18	~ 37
Marginal fallback (silica, Δx=0.5 μm)	1.0×10^{-15}	2.2 s	0.64	~ 190
Heavy levitated projected (Δx=1 μm)	1.0×10^{-14}	53 ms	~ 0	25
MAQRO baseline ($\Delta x \gtrsim 2R$)	1.6×10^{-17}	623 s	0.998	$\sim 10^7$
Levitated, current (Δx=10 nm)	1.0×10^{-18}	4×10^6 s	1.000	3×10^{14}
BMV/QGEM (Δx=250 μm)	1.0×10^{-14}	12 ms	~ 0	25

Table 1. *The G^1-versus-G^2 falsifiability map (coefficient $C = 1$; at the floor $C = 2/\pi$ all decoherence times lengthen by $\pi/2$ and $M_\star$ shifts by $\sim 1.3\times$). The clean discriminant is a picogram-scale single-mass visibility measurement at cryogenic extreme-high vacuum: the envelope target gives a visibility drop to 0.18, decided in a few tens of runs. Sub-threshold masses ($M < M_\star$)—including MAQRO's baseline particle ($\sim 10^{-17}$ kg), which needs a ~ 100 s hold to reach even a marginal $\sim 10^3$-run test—are out of reach or marginal; large-separation (Bose–Marletto–Vedral) platforms self-decohere in milliseconds (the "swamping" regime). All values follow from the exact sphere self-energy and the physical constants.*

- **A locked structural ratio** $a_0/(c\,\Gamma_{\text{total}}) = 2/g(1)$ linking the MOND scale to the cosmic decoherence rate (Section 5): unique to QGC as an internal architecture constraint, but *untestable*, since the cosmic rate is unmeasurably small.

5 The cosmological sector

The same de Sitter thermodynamics that underlies the decoherence programme has consequences on cosmological scales, developed in dedicated papers. This section summarizes the results and their status and gives the single structural relation that ties the sector to the laboratory one.

5.1 The MOND acceleration scale

Treating the entanglement of the de Sitter horizon as an elastic medium yields a covariant field equation $G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G(T^{\text{matter}}_{\mu\nu} + T^{\text{elastic}}_{\mu\nu})$ with an elastic modulus $\kappa = c^4/(8\pi G a_0)$ set by a single acceleration scale. That scale is fixed, with no free parameters, by the de Sitter temperature:

$$a_0 \;=\; \frac{cH_0}{2\pi} \;\approx\; 1.08 \times 10^{-10}\ \text{m/s}^2, \tag{12}$$

which reproduces the empirical MOND scale [30, 7]. The recovery of MOND *phenomenology* from de Sitter thermodynamics is derived; the interpolation function $\mu(x)$ that governs the transition between the Newtonian and deep-MOND regimes is *not* uniquely fixed by the framework—only its limiting behaviours ($\mu \to 1$ at high acceleration, $\mu \to x$ at low) are derived, with the transition region observationally determined. A residual factor-~ 2 tension in galaxy-cluster masses, shared with MOND generally, remains open.

5.2 Holographic dark energy

The same horizon thermodynamics gives a dark-energy density

$$\rho_{\text{DE}} \;=\; \alpha\,\frac{c^2 H^2}{G}, \qquad \alpha \approx 0.082, \tag{13}$$

derived from the holographic framework [6], with the coefficient α understood as the product of a horizon-information ceiling and an epoch factor. The acoustic-peak structure of the cosmic

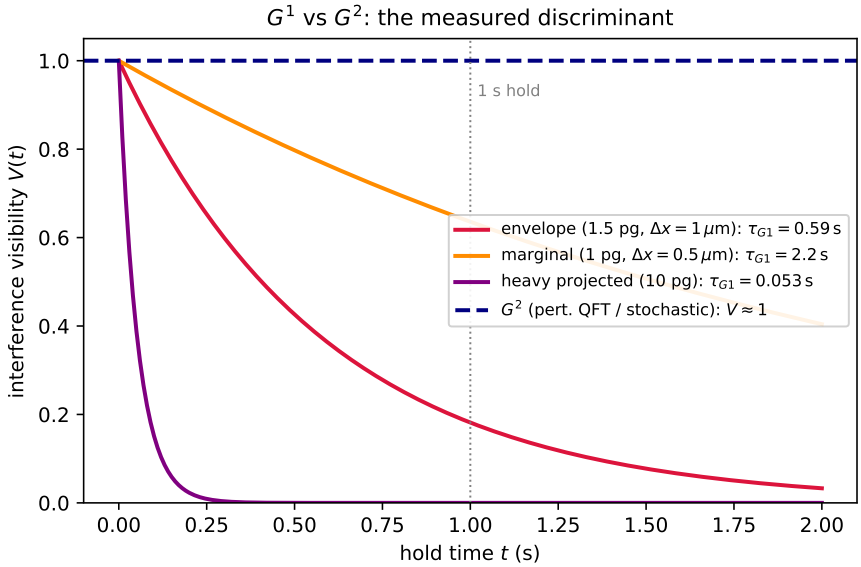

Figure 2. *What an experiment measures: the interference visibility $V(t) = e^{-t/\tau_{\rm dec}}$ for the platform ladder, against the G^2 prediction (perturbative quantum field theory or stochastic gravity), for which $V \approx 1$ over any laboratory hold. The envelope target falls to $V = 0.18$ at one second; the standard prediction is a flat line. The two are separated by the full height of the plot.*

microwave background constrains the early-time behaviour: in the saturated ($w = -1$) realization the early dark-energy fraction is negligible and the peaks coincide with ΛCDM (a consistency check), while a naive constant-fraction realization is excluded by current data.

5.3 Cosmic decoherence

Applying the laboratory decoherence rate to cosmological mass scales naively gives an absurd result—a rate of order 10^{103} Hz—which is a Minkowski-formula artifact. The de Sitter recomputation [5], together with a quantum-cosmology selection of the relevant mass granularity and an additivity bound, gives instead

$$\Gamma_{\rm total} \;=\; g(1)\,\frac{H}{4\pi} \;\approx\; \frac{H}{10\pi} \;\approx\; 7\times10^{-20}\ \text{Hz}, \qquad g(1) = 1 - \tfrac{2}{\pi}\,\mathrm{Si}(1) \approx 0.398, \tag{14}$$

which is sub-causal and physically harmless. The engine of the resolution is the exact identity $\varepsilon_{\rm dS} = G\hbar H^2/(2\pi c^5) = 1/(2S_{\rm dS})$ relating the de Sitter post-Newtonian parameter to the horizon entropy. The result holds as a controlled approximation, with the $\mathcal{O}(G^2)$ residues disclosed. The one premise the additivity bound once carried—that the decoherent-histories readout occurs at horizon crossing rather than at a fine-tuned sub-horizon time—is now derived from standard de Sitter classicalization [31]: a mode registers its which-path phase only when it freezes and squeezes at horizon crossing, while sub-horizon modes oscillate and average away.

5.4 The structural link to the laboratory

The cosmological and laboratory sectors are not independent: they share a single de Sitter frequency, the "horizon clock" [8]. The MOND scale and the cosmic decoherence rate are locked by

$$\frac{a_0}{c\,\Gamma_{\text{total}}} = \frac{2}{g(1)}, \tag{15}$$

a relation with no free parameters: the same 2π of de Sitter thermality appears in both. This is a genuine internal-consistency constraint—the programme has one cosmic input where a phenomenological model would have two independent knobs—but it is *not* a falsifiable prediction, because Γ_{total} is unmeasurably small. It belongs among the programme's structural results, not among its tests.

6 Relation to other approaches

Gravitational decoherence is a crowded field. QGC is positioned here against the three families of proposals it most resembles—where it agrees, where it differs, and what would tell them apart.

6.1 Spontaneous-collapse models (GRW, CSL)

Continuous spontaneous localization and its relatives [11, 12] modify the Schrödinger equation by adding a stochastic, nonlinear localization term with a postulated noise strength and correlation length. They predict decoherence of macroscopic superpositions and are falsifiable, and current experiments already constrain their parameters. QGC differs in kind: it adds no stochastic term and no new fundamental constant. The decoherence is a consequence of the gravitational constraint acting on an otherwise unitary theory, and its rate is fixed by G, $\hbar$, and the configuration through E_{G}, with the only freedom the $\mathcal{O}(1)$ coefficient C of (9). Operationally, collapse models predict a noise spectrum (heating, spontaneous radiation) that QGC does not; conversely, QGC's mass scaling $\Gamma_{\text{dec}} \propto M^2/d$ is sharper than the collapse models' free-parameter freedom. The cleanest discriminator from CSL is that QGC's rate has no adjustable strength: it is the Diósi–Penrose value or it is wrong.

6.2 The Diósi–Penrose proposal

QGC is closest to the Diósi–Penrose (DP) gravitational-decoherence proposal [9, 10], and recovers its rate as the laboratory limit: the decoherence time (1) is the DP form, and the natural coefficient $C = 1$ is the DP value. The difference is one of foundation rather than of phenomenology. DP posits the rate from a gravitational self-energy uncertainty argument; QGC attempts to *derive* it from the constrained modular dynamics (Section 3), and in doing so commits to the object-size (continuous-density) cutoff that lets the rate evade the nucleon-scale spontaneous-radiation bound that constrains the original DP prescription. The two are therefore experimentally near-degenerate in the laboratory regime—both predict the visibility decay of Table 1—but QGC carries a falsifiable structural commitment (the cutoff, the G^1 scaling from the constraint) that DP, as a phenomenological rate, does not.

6.3 Stochastic and semiclassical gravity

Stochastic- and semiclassical-gravity approaches [13] treat the metric as a classical field driven by the quantum stress-energy fluctuations, yielding decoherence from the induced metric noise. This is a G^2 effect—the rate is quadratic in the coupling—and so coincides, at the level of scaling, with perturbative quantum field theory rather than with QGC's G^1 claim. The disagreement is exactly the one the experiments of Section 4 are designed to resolve: stochastic gravity and perturbative QFT predict no laboratory decoherence of a microgram superposition, while QGC predicts nanosecond-scale decoherence. This is the sharpest available test of whether gravity must be quantized: QGC's G^1 rate is a signature of a genuinely quantum (constraint-bearing) gravitational sector, whereas the G^2 rate is what a classical-channel or perturbative-graviton treatment gives.

6.4 Programmes of quantum gravity

QGC is not a candidate fundamental theory of quantum gravity in the sense of loop quantum gravity, string theory, or asymptotic safety; it is a correspondence principle with a near-term prediction. This distinction sets what QGC is and is not accountable for.

Loop quantum gravity and related canonical programmes quantize the gravitational field directly, building a Hilbert space of spin-network states with discrete geometric spectra. QGC takes the opposite stance: it does not quantize the metric, treating geometry instead as an emergent, entropic description dual to the matter sector (Axioms I–II). Where loop quantum gravity predicts Planck-scale discreteness of areas and volumes, QGC treats a minimum length only as a speculative ultraviolet extension on the generalized-uncertainty route, and makes no claim about a fundamental spacetime microstructure. The programmes are not competitors so much as answers to different questions—"what are the quantum degrees of freedom of geometry?" versus "what does treating geometry as emergent predict for a laboratory superposition?"—and QGC's content is the latter.

String theory embeds gravity in a unitary quantum theory with a definite ultraviolet completion, and its holographic incarnation (AdS/CFT) is the deepest existing realization of the geometry-as-entanglement idea that QGC takes as a working principle. QGC borrows the modular and crossed-product technology that AdS/CFT and the de Sitter algebra programme have developed (Section 3), and is consistent with the holographic dictionary, but it does not require a string-scale completion or extra dimensions: its predictions live at the linearized, low-energy level where the gravitational constraint, not the ultraviolet completion, is what matters. Given AdS/CFT, QGC is an attempt to extract a flat-space, laboratory-scale consequence of the same entanglement–geometry correspondence.

Asymptotic safety posits a non-trivial ultraviolet fixed point that renders quantized gravity predictive, and shares with QGC the use of thermodynamic and renormalization-group arguments. It differs in retaining a quantized metric with a running gravitational coupling; QGC keeps G fixed and locates the new physics in the constraint structure rather than in the flow. Of the three, asymptotic safety is the one whose low-energy phenomenology could in principle overlap QGC's, but the two make their sharpest statements in different regimes (the ultraviolet fixed point versus the laboratory decoherence rate), and no direct conflict has been identified.

The honest position is that QGC neither subsumes nor is subsumed by these programmes. It is a correspondence that, if its G^1 rate is confirmed, would constrain any fundamental theory to reproduce that rate; if falsified, it would leave the fundamental programmes untouched, since none of them predicts G^1 decoherence. Its value is the prediction, not a claim to fundamentality.

Approach	Rate scaling	Free strength?	Lab decoherence ($1\,\mu$g, $1\,$mm)
QGC	G^1	no (coeff. $C \in [2/\pi, 1]$)	$\tau_{\rm dec} \approx 1.6\,$ns
Diósi–Penrose	G^1	no	$\tau_{\rm dec} \approx 1.6\,$ns
CSL/GRW	— (postulated)	yes	parameter-dependent
Stochastic gravity	G^2	no	$\sim 10^{18}\,$yr
Perturbative QFT	G^2	no	$\sim 10^{18}\,$yr

Table 2. *Discriminators among gravitational-decoherence approaches. QGC and Diósi–Penrose are laboratory-degenerate (G^1, no free strength); the G^1-versus-G^2 split is the decisive experimental question, and CSL is separated from both by its adjustable noise strength and its spontaneous-radiation signature.*

6.5 Summary of the discriminators

7 Open problems and roadmap

The problems that remain are ordered by how much their resolution would move the programme.

7.1 The decisive problem: closing the finite-diamond construction

The single problem whose resolution would change the programme's standing is the closure of the finite-diamond construction at full rigour—equivalently, proving the operator identity (7) and pinning the coefficient (9) to a single value. By Section 3, the construction reduces to one well-posed question: the clock-model normalization that fixes C. The trace fixes the modular temperature and the metric; what remains is to show that the fringe-visibility (Bures) diffusion coefficient is *exactly* $E_{\rm G}/\hbar$ rather than an $\mathcal{O}(1)$ multiple, which requires controlling the spectral shape of the constraint-borne noise. A proof here would simultaneously legitimize the operator identity, the rate extraction, and the $C = 1$ selection, and would turn the central claim from conjecture into theorem. A demonstration that the construction is ill-posed, or an experiment showing G^2, would refute it. This is the fulcrum of the whole programme.

7.2 Theory problems of secondary impact

- **The Bose–Marletto–Vedral rate inconsistency** remains to be resolved before that prediction is firm: the arm-set and radius-saturated self-decoherence rates disagree by $\sim 10^2$ at experimental parameters, and the visibility law is contingent on the choice.
- **Rigorous $\mathcal{O}(G^2)$ remainder bounds** for the coherent-state argument exist only at the level of power counting; a genuine operator bound is open.
- **The MOND interpolation function** $\mu(x)$ is not derivable from the current entropy-equilibrium argument; only its limits are fixed. The galaxy-cluster factor-~ 2 tension is the associated phenomenological gap.
- **The edge-mode and trace-anomaly bookkeeping** of the finite-diamond construction, while shown not to affect the rate coefficient, remains to be computed for a complete entropy ledger.

7.3 The experimental roadmap

The programme's fate is ultimately experimental. The decisive measurement is a picogram-scale solid-object superposition (just above the threshold $M_\star \approx 0.76$ pg), held for of order a second, in cryogenic extreme-high vacuum, with environmental decoherence budgeted below the G^1 rate (Section 4). The required mass is some three orders of magnitude above today's ground-state-cooled levitated spheres, and the central challenge—a micrometre-scale spatial superposition of a sub-micrometre sphere—is at the frontier of the field and is the goal of MAQRO-class proposals on a roughly 2028–2035 horizon (with the caveat that MAQRO's nominal $\sim 10^{-17}$ kg particle is itself sub-threshold and would need a heavier sphere). A positive result would establish G^1 and with it the quantum character of the gravitational sector; a clean null below the predicted timescale would falsify the programme's central claim. Either outcome is informative, which is the property a research programme should want of its predictions. In parallel, wide-binary statistics (Gaia DR4 and beyond) and high-redshift rotation curves test the cosmological MOND scale (12) and its possible redshift dependence $a_0(z) = cH(z)/(2\pi)$.

7.4 What would falsify the programme

The programme is falsified by: (i) an experiment in the decidable window (Table 1) that maintains coherence well past τ_{dec}, contradicting G^1; (ii) a demonstration that the finite-diamond construction is ill-posed, removing the basis for the operator identity; or (iii) a cosmological measurement of $a_0(z)$ inconsistent with $cH(z)/(2\pi)$. None of these is currently realized, and the first is within experimental reach this decade.

8 Conclusion

The Quantum-Geometric Correspondence treats quantum mechanics and general relativity as complementary projections of a single quantum-informational reality, and asks what that stance predicts. The answer divides cleanly. The kinematic and equilibrium layer—three primitive axioms recovering the Bekenstein–Hawking entropy, the holographic bound, the Unruh temperature, the Einstein equations, and a thermodynamic dark-energy density—is on standard ground. The single new dynamical prediction, gravitational decoherence at the G^1 rate $\Gamma_{\text{dec}} = CE_{\text{G}}/\hbar$, separates into an established energy scale $E_{\text{G}} = GM^2/d$ and a conjectured operator identity $K = 2\pi H_{\text{phys}}$ that converts the scale into a rate. The latter has been reduced, through the finite-diamond construction, to a single well-posed open question—the clock-model normalization that fixes the coefficient $C \in [2/\pi, 1]$—with the energy scale, the geometric flow (now including the physical graviton), the intrinsic observer clock, the decoupling of the entropy residues, and the infrared safety all established along the way.

The balance is honest: the physics is intact and the falsifiable claim is unchanged, while the central identity stands as a strong conjecture rather than a theorem, and no part of the programme has yet been confirmed experimentally. The decisive virtue of the prediction is that it is testable: a picogram superposition held for a second in cryogenic vacuum distinguishes the G^1 rate from the G^2 rate of perturbative quantum field theory and stochastic gravity by some thirty-four orders of magnitude, in a few tens of interferometric runs, with an experiment at the frontier of—but not beyond—current levitated-optomechanics capability. Whether spacetime must be quantized, and whether it decoheres superpositions on the Diósi–Penrose timescale, are questions this programme renders sharp and an experiment this decade can answer. That is the outlook: a conservative correspondence with one open construction and one falsifiable number, waiting on a measurement.

Appendices

A The rate coefficient: derivation and the residual freedom

This appendix recapitulates how the coefficient C in $\Gamma_{\rm dec} = C\, E_{\rm G}/\hbar$ is bounded, why it is not yet uniquely pinned, and what the open question is, distinguishing throughout a bound from a derivation.

A.1 The free field saturates: an exponent, not a rate

The free-field decoherence exponent for the two-branch coherent-state overlap is

$$\Gamma(t) \;=\; \int \frac{d^3k}{(2\pi)^3}\, |\delta\alpha(k)|^2\, 2(1-\cos\omega_k t), \qquad \omega_k = c|k|, \tag{16}$$

with the equilibrium amplitude $\delta\alpha(k) \propto (GM/k^2)$ set by the Newtonian fields of the two branches. By the Riemann–Lebesgue lemma the oscillatory term averages away, and the exponent *saturates* at a static value $\Gamma_{\rm sat} = \xi \ln(d/\ell_{\rm P})$ with $\xi = GM^2/(\hbar c)$. This is a dimensionless coherence-loss exponent—a one-time overlap reduction $|\langle\Phi_L|\Phi_R\rangle|^2 = e^{-\Gamma_{\rm sat}}$—not a rate. The free graviton bath is super-Ohmic, with spectral density $J(\omega) \propto \omega^3$ and $J(0^+) = 0$, so it carries no zero-frequency weight and produces no sustained (Markovian) decoherence on its own. The energy scale is present in (16); the rate is not.

A.2 The constraint supplies the rate

The linear-in-time rate is not a free-field effect: it is enforced by the Wheeler–DeWitt Hamiltonian constraint, which ties the gravitational field to its source and prevents the branches from sharing a single un-sourced field configuration. In the constrained influence functional the static overlap is converted into a sustained dephasing, and the natural Markovian rate is

$$\frac{d\rho_{\rm matter}}{dt} \;=\; -\frac{i}{\hbar}[H_{\rm eff},\rho] - \frac{\Gamma}{2}\,[X,[X,\rho]], \qquad \Gamma = \frac{E_{\rm G}}{\hbar}, \tag{17}$$

i.e. $C = 1$, "up to an $\mathcal{O}(1)$ coefficient that depends on the detailed clock model." That residual $\mathcal{O}(1)$ is the whole of the open question.

A.3 What the trace fixes, and the two-observable residue

In the finite-diamond (Type II) formulation the canonical trace removes the normalization freedom that earlier gave a factor-four window: it fixes the modular temperature $\beta_{\rm mod} = 2\pi$, the Kubo–Martin–Schwinger variance of the modular Hamiltonian, and hence the Bures–Fisher metric on the two-branch state space. What survives is not a normalization ambiguity but a genuine distinction between two operationally different observables built from the same trace-fixed metric:

$$\text{continuous fringe-visibility (Bures) rate:} \quad \Gamma = E_{\rm G}/\hbar \;\Rightarrow\; C = 1, \tag{18}$$

$$\text{Margolus–Levitin orthogonalization speed:} \quad \Gamma_\perp = (2/\pi)\, E_{\rm G}/\hbar \;\Rightarrow\; C = 2/\pi. \tag{19}$$

Their ratio $C_{\text{vis}}/C_{\perp} = \pi/2$ is exactly the width of the window (9); the window is the ratio of two different observables, not the spread of one. An interferometer measures fringe-visibility decay, so the trace-fixed metric *selects* $C = 1$ for the falsifiable prediction, and the $2/\pi$ anchor is a distinct quantity (a speed-limit ceiling), not a competing value.

A.4 The honest residue

What the trace does *not* supply is the spectral shape that would make the Bures-diffusion coefficient *exactly* $E_{\mathrm{G}}/\hbar$ rather than $\kappa\, E_{\mathrm{G}}/\hbar$ with $\kappa = \mathcal{O}(1)$. The Kubo–Martin–Schwinger analyticity strip gives a correlation time $\tau_c \sim \beta_{\text{mod}}$ and hence $\kappa = \mathcal{O}(1)$, Markovian to leading order; but the super-Ohmic spectrum offers no zero-frequency knob to set $\kappa = 1$ on the nose, so the exactness of $C = 1$ is constraint-borne and is precisely the content of the operator equation $K = 2\pi H_{\text{phys}}$. In summary: the bound $C \in [2/\pi, 1]$ is derived; $C = 1$ is selected for the measured observable by the trace-fixed metric; and the uniqueness of $C = 1$ is conjectured, equivalent to closing the finite-diamond construction at full rigour (Appendix B).

B Status of the finite-diamond construction

Section 3 summarizes the finite-diamond construction in prose; the table below sets out what is established and what remains open, so that the single open lever is visible at a glance. Entries marked as symbolic or lattice rest on the corresponding computation.

The structural content is that the diamond modular Hamiltonian is a single object whose *gap* is the energy scale,

$$\frac{\hbar}{2\pi}\,\Delta K = E_{\mathrm{G}} \quad \text{(normalization-independent)}, \tag{20}$$

and whose *flow* on the intrinsic clock is the rate $\Gamma_{\text{dec}} = C\, E_{\mathrm{G}}/\hbar$, with the infrared-safe form $\Gamma_{\text{dec}} = 2N(1 - a^2)\, E_{\mathrm{G}}/\hbar$ for a source at fractional radius a inside a diamond of any size, N the clock normalization. The energy scale is fixed; the coefficient $C = 2N$ (for a centered source) is the one quantity the construction does not yet pin, and pinning it is equivalent to proving the operator equation.

C The exact sphere self-energy and the falsification numerics

The decoherence times of Section 4 use the exact Diósi–Penrose self-energy of a uniform sphere, not the point-mass form. For a sphere of mass M and radius $R = (3M/4\pi\rho)^{1/3}$ in a superposition of separation Δx, the gravitational self-energy difference between the branches is

$$E_{\mathrm{G}}(\Delta x) \;=\; \frac{GM^2}{R} \times \begin{cases} \frac{1}{2}s^2 - \frac{3}{16}s^3 + \frac{1}{160}s^5, & s = \Delta x/R \le 2, \\ \frac{6}{5} - R/\Delta x, & \Delta x \ge 2R, \end{cases} \tag{21}$$

with the two limiting behaviours

$$E_{\mathrm{G}} \to \tfrac{1}{2}\,\frac{GM^2}{R^3}\,\Delta x^2 \quad (\Delta x \ll R,\ \text{tidal}), \qquad E_{\mathrm{G}} \to 1.2\,\frac{GM^2}{R} \quad (\Delta x \gtrsim 2R,\ \text{saturated}). \tag{22}$$

The decoherence time is $\tau_{\rm dec} = \hbar/E_{\rm G}$, the visibility $V(t) = e^{-t/\tau_{\rm dec}}$.

The threshold. In the saturated regime $E_{\rm G} = 1.2\,GM^2/R$ with $R \propto M^{1/3}$, so $\tau_{\rm dec} \propto M^{-5/3}$: there is a sharp threshold mass below which the superposition is undecidable on laboratory timescales. Setting $\tau_{\rm dec} = 1$ s gives, for silica ($\rho = 2200$ kg/m^3),

$$M_\star \;=\; \left[\frac{\hbar}{1.2\,G}\left(\frac{3}{4\pi\rho}\right)^{1/3}\right]^{3/5} \;\approx\; 7.6\times10^{-16}\ \text{kg} \quad (R\approx 0.44\,\mu\text{m}). \tag{23}$$

The 5σ run count. A visibility deficit $1-V$ is resolved against the standard null $V=1$ with a per-run shot-noise error $\sigma_V \approx 1/\sqrt{N}$, so 5σ discrimination requires $(1-V)\sqrt{N} \geq 5$, i.e.

$$N_{5\sigma} \;=\; \frac{25}{(1-V)^2}. \tag{24}$$

This is an idealized floor for unit per-run fringe contrast; a real per-run contrast $c<1$ multiplies it by $1/c^2$.

The environmental budget. The test is clean only if residual-gas decoherence $\Gamma_{\rm gas} = n\sigma\bar{v}$ (with $\sigma = \pi R^2$, $n = P/k_{\rm B}T$) is slower than $\Gamma_{\rm dec} = 1/\tau_{\rm dec}$, requiring

$$P \;\lesssim\; \frac{k_{\rm B}T}{\pi R^2\bar{v}\,\tau_{\rm dec}}, \qquad \bar{v} = \sqrt{\frac{8k_{\rm B}T}{\pi m_{\rm gas}}}, \tag{25}$$

which for the envelope target (1.5 pg, $\Delta x = 1\,\mu$m, $T = 4$ K, helium) gives $P \lesssim 7\times10^{-15}$ mbar—cryogenic extreme-high vacuum. Every value in Table 1 and Figures 1–2 follows from (21)–(25) and the physical constants.

References

[1] M. Sperzel, *The Quantum-Geometric Correspondence: Three Axioms and Gravitational Decoherence at Order G*, QGC programme paper, `foundations/canonical-theory` (2026).

[2] M. Sperzel, *Geometric Modular Flow for Gravitons in a Causal Diamond: Dimensional Reduction and the Type II Crossed Product for a Non-Conformal Field*, QGC programme paper, `foundations/graviton-modular-flow` (2026).

[3] M. Sperzel, *A Mass-Independent Visibility Suppression in the Bose–Marletto–Vedral Experiment from Quantum-Geometric Correspondence*, QGC programme paper, `decoherence/bmv-prediction` (2026).

[4] M. Sperzel, *The Gravitational Decoherence Floor: Platform Reach, Finite-Size Scaling Laws, and the Entanglement-Decay Signature*, QGC programme paper, `decoherence/decoherence-signatures` (2026).

[5] M. Sperzel, *Cosmic Decoherence on de Sitter: Resolving the Naive Infinity*, QGC programme paper, `decoherence/cosmic-decoherence` (2026).

[6] M. Sperzel, *Holographic Cosmology: Dark Energy and the Arrow of Time from One Horizon Entropy Budget*, QGC programme paper, `cosmology/holographic-cosmology` (2026).

[7] M. Sperzel, *Emergent Gravity from Entanglement Equilibrium: MOND Phenomenology and Covariant Extension*, QGC programme paper, `cosmology/emergent-gravity-mond` (2026).

[8] M. Sperzel, *The Horizon Clock: A Single de Sitter Frequency Behind the MOND Scale, Cosmic Decoherence, and the Chaos Bound*, QGC programme paper, `cosmology/horizon-clock` (2026).

[9] L. Diósi, *A universal master equation for the gravitational violation of quantum mechanics*, Phys. Lett. A **120**, 377 (1987). doi:10.1016/0375-9601(87)90681-5

[10] R. Penrose, *On gravity's role in quantum state reduction*, Gen. Relativ. Gravit. **28**, 581 (1996). doi:10.1007/BF02105068

[11] G. C. Ghirardi, A. Rimini, and T. Weber, *Unified dynamics for microscopic and macroscopic systems*, Phys. Rev. D **34**, 470 (1986). doi:10.1103/PhysRevD.34.470

[12] A. Bassi, K. Lochan, S. Satin, T. P. Singh, and H. Ulbricht, *Models of wave-function collapse, underlying theories, and experimental tests*, Rev. Mod. Phys. **85**, 471 (2013). doi:10.1103/RevModPhys.85.471

[13] M. P. Blencowe, *Effective field theory approach to gravitationally induced decoherence*, Phys. Rev. Lett. **111**, 021302 (2013). doi:10.1103/PhysRevLett.111.021302

[14] S. Bose, A. Mazumdar, G. W. Morley, H. Ulbricht, M. Toroš, M. Paternostro, A. A. Geraci, P. F. Barker, M. S. Kim, and G. Milburn, *Spin entanglement witness for quantum gravity*, Phys. Rev. Lett. **119**, 240401 (2017). doi:10.1103/PhysRevLett.119.240401

[15] C. Marletto and V. Vedral, *Gravitationally induced entanglement between two massive particles is sufficient evidence of quantum effects in gravity*, Phys. Rev. Lett. **119**, 240402 (2017). doi:10.1103/PhysRevLett.119.240402

[16] R. Kaltenbaek et al., *Macroscopic quantum resonators (MAQRO): 2015 update*, EPJ Quantum Technol. **3**, 5 (2016). doi:10.1140/epjqt/s40507-016-0043-7

[17] U. Delić, M. Reisenbauer, K. Dare, D. Grass, V. Vuletić, N. Kiesel, and M. Aspelmeyer, *Cooling of a levitated nanoparticle to the motional quantum ground state*, Science **367**, 892 (2020). doi:10.1126/science.aba3993

[18] J. J. Bisognano and E. H. Wichmann, *On the duality condition for quantum fields*, J. Math. Phys. **17**, 303 (1976). doi:10.1063/1.522898

[19] M. Takesaki, *Duality for crossed products and the structure of von Neumann algebras of type III*, Acta Math. **131**, 249 (1973). doi:10.1007/BF02392041

[20] H. Casini, M. Huerta, and R. C. Myers, *Towards a derivation of holographic entanglement entropy*, J. High Energy Phys. **05** (2011) 036. doi:10.1007/JHEP05(2011)036

[21] V. Chandrasekaran, R. Longo, G. Penington, and E. Witten, *An algebra of observables for de Sitter space*, J. High Energy Phys. **02** (2023) 082. doi:10.1007/JHEP02(2023)082

[22] K. Jensen, J. Sorce, and A. J. Speranza, *Generalized entropy for general subregions in quantum gravity*, J. High Energy Phys. **12** (2023) 020. doi:10.1007/JHEP12(2023)020

[23] V. Benedetti and H. Casini, *Entanglement entropy of linearized gravitons in a sphere*, Phys. Rev. D **101**, 045004 (2020). doi:10.1103/PhysRevD.101.045004

[24] M. Huerta and G. van der Velde, *Modular Hamiltonian of the scalar in the semi-infinite line: dimensional reduction for spherically symmetric regions*, J. High Energy Phys. **06** (2023) 097. doi:10.1007/JHEP06(2023)097

[25] T. Regge and J. A. Wheeler, *Stability of a Schwarzschild singularity*, Phys. Rev. **108**, 1063 (1957). doi:10.1103/PhysRev.108.1063

[26] F. J. Zerilli, *Effective potential for even-parity Regge–Wheeler gravitational perturbation equations*, Phys. Rev. Lett. **24**, 737 (1970). doi:10.1103/PhysRevLett.24.737

[27] T. Jacobson, *Thermodynamics of spacetime: the Einstein equation of state*, Phys. Rev. Lett. **75**, 1260 (1995). doi:10.1103/PhysRevLett.75.1260

[28] J. D. Bekenstein, *Black holes and entropy*, Phys. Rev. D **7**, 2333 (1973). doi:10.1103/PhysRevD.7.2333

[29] W. G. Unruh, *Notes on black-hole evaporation*, Phys. Rev. D **14**, 870 (1976). doi:10.1103/PhysRevD.14.870

[30] M. Milgrom, *A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis*, Astrophys. J. **270**, 365 (1983). doi:10.1086/161130

[31] D. Polarski and A. A. Starobinsky, *Semiclassicality and decoherence of cosmological perturbations*, Class. Quantum Grav. **13**, 377 (1996). doi:10.1088/0264-9381/13/3/006

Ingredient of the construction	Status
Flat-space Type III_1 obstruction is real (no trace; cross-wedge support)	established
Reduction of the frontier to the diamond crossed product	derived
Type II_∞ algebra + canonical trace for $\mathcal{A}(D) \rtimes_\sigma \mathbb{R}$	established (Takesaki)
Diamond conformal-Killing flow = exact $\mathbb{R} \times \mathbb{H}^3$ Killing flow	established (symbolic)
Geometric flow for the physical graviton (tower reduction)	derived
Regge–Wheeler/Zerilli $\to$ scalar degeneration at $M \to 0$	established (symbolic + lattice)
Intrinsic observer worldline; uniform clock; $\beta_{\mathrm{mod}} = 2\pi$	established (exact)
Bounded-below observer Hamiltonian (II_1 refinement)	scoped (S_0, C-irrelevant)
Edge-mode / center sector at ∂B_R	scoped (S_0, C-irrelevant)
Four-dimensional trace anomaly	scoped (S_0, C-irrelevant)
Modular flow invariant under $K \to K + c\mathbb{1}$ (the decoupling)	established (exact)
Rate infrared-safe (independent of the diamond size R)	established (exact)
Hamiltonian constraint = modular generator; imposing it selects $K = 2\pi H_{\mathrm{phys}}$	derived (solvable model)
Diamond nuclearity $\Rightarrow$ split property (type-I approximant for the dressing)	established (spectral gap + cited theorem)
Operator-equation error = exterior dressing weight, $(d/h)^4$-suppressed for the physical source	derived
Assembled bound $\\|K - 2\pi H_{\mathrm{phys}}\\| \le C\,(d/h)^4 + O(G^2)$	derived (Gaussian sector)
Split constant C: finite under interactions; explicit free-field value	derived (controlled model)
Bi-local residual: onsets at $O(G^2)$; closed-form kernel value, norm bound with constant 1	derived (controlled model)
Residual size at the laboratory benchmark ($\lesssim 3 \times 10^{-64}$ of the leading rate)	derived (controlled model)
Composition of all the above (one geometry)	established
Exact coefficient C **(clock-model normalization)**	**conjectured**, $C \in [2/\pi, 1]$
Operator equation $K = 2\pi H_{\mathrm{phys}}$ (interacting values of C and the residual)	**conjectured** (holds in expectation values)

Table 3. *What is established and what remains open in the finite-diamond construction. Everything above the rule is proven, derived, or scoped (shown not to affect the rate coefficient); the single open lever is the clock-model normalization that fixes C, dual to the operator equation. The three scoped residues are entropy-constant (S_0) effects, decoupled from C by the exact invariance of the modular flow under an additive shift. The two quantities below the rule now carry explicit controlled-model values (free-field split constant; closed-form bi-local kernel bound); what remains conjectural is their values in the interacting theory — which is precisely what the G^1-versus-G^2 measurement probes.*