Cosmic Gravitational Decoherence on a de Sitter Background:

Resolving a Causality Paradox in Quantum-Geometric Duality

Paper N of the Quantum-Geometric Duality Series

Marc Sperzel*

Independent Researcher

MSci Physics, King's College London

May 22, 2026

Abstract

The gravitational decoherence rate $\Gamma = G\,\Delta m^2/(\hbar d)$ of Quantum-Geometric Duality (QGD), applied at cosmic scale with Δm the mass of the observable universe and d the Hubble radius R_{H}, returns $\Gamma \sim 10^{103}$ Hz—exceeding the causal bound $c/R_{\mathrm{H}} \sim H_0$ by 121 orders of magnitude. This internal contradiction, logged within the framework as weakness "W2," is the subject of this paper. We show that the divergence is an artifact of using the Minkowski graviton propagator at separations $d \sim R_{\mathrm{H}}$, where it is not valid. Recomputed on a de Sitter background with Bunch–Davies modes, the coherent-state decoherence functional acquires a physical infrared cutoff at the horizon scale: super-horizon modes freeze and cease to drive decoherence. The rate gains a closed-form form factor, $\Gamma_{\mathrm{dS}} = (G\,\Delta m^2/\hbar d)\,g(Hd/c)\,t$ with $g(x) = 1 - (2/\pi)\,\mathrm{Si}(x)$, which recovers the Minkowski result for $Hd/c \ll 1$ and is finite and causal at $d = R_{\mathrm{H}}$. The cosmic rate is governed not by the total mass M_U but by the elementary branch granularity of the cosmic wavefunction; arguments from quantum cosmology identify this granularity with the de Sitter thermal mass $\Delta m_{\mathrm{dS}} = \hbar H/(2\pi c^2)$ per horizon mode. Summed over the S_{dS} holographic horizon modes, the cosmic decoherence rate is $\Gamma_{\mathrm{total}} = g(1)\,H/(4\pi) \approx H/(10\pi) \approx 7 \times 10^{-20}$ Hz, manifestly sub-causal. An internal-consistency program addresses the residual approximations—linear additivity, the spin-2 graviton, gauge-invariant rate extraction, and all-orders secular control—leaving a small set of explicitly stated premises. The result resolves an inconsistency internal to QGD; we state honestly which premises it rests on and what falsifiable consequences it carries.

Contents

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

1 Introduction

A physical theory is only as trustworthy as its weakest internal contradiction. Quantum-Geometric Duality (QGD)—the proposal that quantum mechanics and general relativity are dual descriptions of a single informational substrate—predicts that a spatial superposition of a mass loses coherence through its gravitational field at a rate

$$\Gamma = \frac{G\,\Delta m^2}{\hbar\, d}, \tag{1}$$

where Δm is the mass difference between the two branches and d their spatial separation [1, 2]. The rate (1) has the same parametric form as the gravitational decoherence rate of the Diósi–Penrose conjecture [28, 27], to which QGD lends a specific microphysical mechanism. For laboratory parameters it is benign: a microgram particle separated by a millimetre decoheres in roughly a nanosecond. Equation (1) was derived in linearized gravity by imposing the Wheeler–DeWitt constraint on the Feynman–Vernon influence functional, and its laboratory consequences are the subject of Papers A and K of this series.

The trouble appears when Eq. (1) is pushed to its largest conceivable scale. Insert the mass of the observable universe, $M_U \sim 1.5 \times 10^{53}$ kg, for Δm, and the Hubble radius $R_{\mathrm{H}} = c/H_0 \sim 1.3 \times 10^{26}$ m for d. The formula returns

$$\Gamma_{\mathrm{naive}}(M_U, R_{\mathrm{H}}) \sim \frac{c^5}{4\,G\,\hbar\,H_0} \sim 10^{103}\ \mathrm{Hz}. \tag{2}$$

No physical decoherence process can run faster than information can cross the system. For a system the size of the Hubble volume the causal bound on any rate is

$$\Gamma \le \frac{c}{R_{\mathrm{H}}} = H_0 \sim 2.2 \times 10^{-18}\ \mathrm{Hz}. \tag{3}$$

Equation (2) exceeds Eq. (3) by 121 orders of magnitude. Within the QGD framework this discrepancy was recorded as an open weakness, labelled "W2," first flagged in the framework's analysis of cosmic decoherence and the arrow of time [3, 4]. It is a genuine internal contradiction: a formula the framework relies on, applied within the framework's own stated domain, returns an answer the framework's own causality requirements forbid. Until W2 is resolved, the cosmic-scale behaviour of QGD gravitational decoherence cannot be regarded as understood.

1.1 Why the naive extrapolation is suspect

There are two distinct ways a formula like Eq. (1) can fail at a new scale. It can be *wrong*—resting on physics that does not hold there—or the *inputs* can be wrong—Δm and d not being the quantities the formula was meant to take. We will argue that both failures occur in the naive extrapolation, and that diagnosing them separately is what resolves W2.

The first failure is a matter of the propagator. Equation (1) is built from the graviton two-point function on a *Minkowski* background. That two-point function has no infrared scale: arbitrarily long-wavelength modes contribute to the mode integral, and on flat space this is mathematically legitimate after self-energy renormalization. But the actual universe is not Minkowski at separations approaching R_{H}; it is, to an excellent approximation, de Sitter, with a Hubble scale that provides exactly the infrared scale Minkowski lacks. Using the flat-space propagator at $d \sim R_{\rm H}$ is using the wrong Green's function in precisely the regime where the difference matters most.

The second failure concerns the inputs. The cosmic wavefunction does not jump between a configuration with all of M_U *here* and the same mass an entire Hubble radius *there*. Inserting $\Delta m = M_U$ computes the decoherence rate between two maximally distant branches—a quantity as physically meaningless as the decoherence rate one would obtain by coherently displacing every molecule of a gas by one metre. The physically relevant question is the rate at which the cosmic wavefunction loses coherence through its *elementary* branch steps, and the elementary step is not M_U.

1.2 What this paper does

This paper resolves W2. The resolution has a structural part and a quantitative part, and both are needed.

The structural part (Section 2) is the de Sitter recomputation. We redo the coherent-state decoherence functional of Paper K on a fixed de Sitter background, with the linearized graviton in the Bunch–Davies vacuum. The de Sitter mode functions automatically suppress modes with comoving wavenumber below $k_H = aH/c$: such modes are super-horizon, they freeze rather than oscillate, and frozen modes do not generate time-dependent decoherence. This is the infrared regulator the Minkowski calculation was missing. The decoherence exponent factorizes as

$$\Gamma_{\rm dS}(t, d, H) = \frac{G\,\Delta m^2}{\hbar\, d}\, g\!\left(\frac{Hd}{c}\right) t, \qquad g(x) = 1 - \frac{2}{\pi}\,\mathrm{Si}(x), \tag{4}$$

where Si is the sine integral. The form factor g recovers the Minkowski result, $g(0) = 1$, for $Hd/c \ll 1$; it is finite at the horizon, $g(1) = 1 - (2/\pi)\,\mathrm{Si}(1) = 0.398$; and it vanishes for super-horizon separations, $g(x \gg 1) \to 0$. The catastrophic divergence of Eq. (2) is gone: at $d = R_{\rm H}$ the rate is finite and proportional to H.

The quantitative part (Section 3) supplies the input the structural part leaves open. The de Sitter recomputation gives a rate per unit Δm^2 but does not by itself select Δm; inserting $\Delta m = M_U$ still returns 10^{103} Hz even on de Sitter. We therefore ask the question native to quantum cosmology: in the Wheeler–DeWitt framework, what is the smallest mass-distribution difference separating two physically distinct branches of the cosmic wavefunction? Several independent lines of argument—a thermal-distinguishability floor, the de Sitter horizon first law, the predictability sieve, and a thermality-free route through the Wheeler–DeWitt branch structure—converge on the de Sitter thermal mass,

$$\Delta m_{\rm dS} = \frac{k_B T_{\rm dS}}{c^2} = \frac{\hbar H}{2\pi c^2}, \tag{5}$$

as the branch granularity per horizon mode. Summing the per-mode rate over the $S_{\rm dS} = \pi c^5/(G\hbar H^2)$ holographic horizon modes gives the cosmic decoherence rate

$$\Gamma_{\rm total} = g(1)\,\frac{H}{4\pi} \approx \frac{H}{10\pi} \approx 7 \times 10^{-20}\ \mathrm{Hz}, \tag{6}$$

finite, proportional to H, and comfortably below the causal bound (3).

Section 4 addresses the approximations on which Eqs. (4)–(6) rest. The result was developed as a controlled approximation and then subjected to an internal-consistency program: the linear additivity of per-mode rates, the use of a scalar proxy for the spin-2 graviton, the gauge dependence of the extracted rate, and the all-orders behaviour of the known de Sitter graviton secular series were each examined in turn. We summarize the outcome of that program—the result survives each check, with stated residuals—without recapitulating it as a sequence of separate episodes.

Section 5 is an honest accounting. The result resolves an inconsistency *internal* to QGD: it does not establish QGD, which remains a framework without direct experimental confirmation, and it rests on premises we state explicitly—chiefly the de Sitter Euclidean periodicity $2\pi/H$ and a mild geometric premise on the quantum-cosmological readout time. We also identify the program-wide assumption the whole framework carries, the identification of a gravitational energy scale with a decoherence rate, which this paper inherits rather than closes. Finally we set out the falsifiable content: what an observation, or a future calculation, would have to show to overturn the picture.

Series context. This paper belongs to a series on gravitational effects at the quantum–classical interface. Paper A presents the laboratory gravitational-decoherence prediction; Paper B develops holographic dark energy; Paper C is the axiomatic framework of QGD; Paper D treats emergent gravity; Paper K derives the G^1 scaling from the constrained influence functional [2]. The present paper extends the Paper A/K decoherence formula to cosmic scale and resolves the causality paradox that extension exposes. It is self-contained but cross-references the others for extended discussion.

2 The de Sitter recomputation

The starting point for resolving W2 is to identify, precisely, where the Minkowski decoherence formula fails. This section does so, then redoes the calculation on the background the cosmic problem actually demands.

2.1 The domain of validity of the Minkowski formula

The Paper K derivation of Eq. (1) represents the gravitational field sourced by a classical mass distribution as a coherent state of the linearized graviton. Two mass configurations differing by Δm at separation d source two coherent states $|\Phi_L\rangle$, $|\Phi_R\rangle$, and the decoherence exponent is the squared norm of their amplitude difference,

$$\Gamma(t) = \left\|\alpha_L(t) - \alpha_R(t)\right\|^2, \tag{7}$$

the norm being the L^2 norm over graviton modes. On a Minkowski background the modes are plane waves and the integral runs over all wavenumbers down to $k = 0$. The infrared end of that integral is rendered finite only by self-energy renormalization—a subtraction that is legitimate on flat space because the divergent piece is k-independent and cancels in the difference $\Gamma_L - \Gamma_R$.

What makes this work on Minkowski is the absence of any infrared scale: there is nothing in flat space to distinguish a mode of wavelength 1 m from one of wavelength 10^{26} m. Both oscillate; both contribute to the time-dependent overlap (7) in the same way. For laboratory separations d this is harmless, because the angular structure of the source already suppresses the longest-wavelength contributions. But as d grows toward the Hubble radius, the calculation begins to draw on modes whose wavelength exceeds the size of the causal patch—and on flat space nothing stops it. The Minkowski formula has no built-in knowledge that such modes are physically special. Its domain of validity is $Hd/c \ll 1$; the naive cosmic extrapolation (2) sits at $Hd/c \sim 1$, outside it.

The actual universe supplies the missing scale. On scales approaching $R_{\rm H}$ the geometry is de Sitter, and de Sitter has a horizon. Modes longer than the horizon are causally frozen. The correct calculation must therefore use the de Sitter graviton propagator, which knows about the horizon, in place of the Minkowski one, which does not.

2.2 Linearized graviton on de Sitter

We work on a fixed de Sitter background in flat (spatially homogeneous) slicing,

$$ds^2 = -c^2\,dt^2 + a(t)^2\,d\mathbf{x}^2, \qquad a(t) = e^{Ht}, \tag{8}$$

or, in conformal time τ defined by $a\,d\tau = c\,dt$ so that $a(\tau) = -1/(H\tau)$ with $\tau \in (-\infty, 0)$,

$$ds^2 = a(\tau)^2\big(-c^2\,d\tau^2 + d\mathbf{x}^2\big). \tag{9}$$

The Hubble radius is $R_{\rm H} = c/H$ and the horizon area $A_H = 4\pi R_{\rm H}^2$. We treat the mass distribution as a classical source for the linearized graviton, neglecting back-reaction of the graviton on the source—the same leading-order approximation used in Paper A.

In transverse-traceless (TT) gauge the linearized Einstein equations on de Sitter reduce, per polarization, to a wave equation. Writing the field as $\Phi_k(\tau) = \chi_k(\tau)/a(\tau)$, the transverse-traceless graviton obeys exactly the mode equation of a massless minimally coupled scalar [7, 6],

$$\chi_k'' + \Big(k^2 - \frac{2}{\tau^2}\Big)\,\chi_k = 0, \tag{10}$$

with primes denoting $d/d\tau$. Throughout the bulk of the derivation we therefore work with a scalar proxy Φ, each TT polarization being treated as such a scalar; that this proxy faithfully reproduces the full spin-2 result, rather than merely approximating it, is established in Section 4.2.

The Bunch–Davies mode functions—selecting positive-frequency behaviour in the far past $\tau \to -\infty$—are

$$\Phi_k(\tau) = \frac{H}{\sqrt{2k^3}}\,(1 + ik\tau)\,e^{-ik\tau}. \tag{11}$$

These mode functions encode the physics that the Minkowski calculation lacked. They split into three classes by the value of $k\tau \propto k_{\rm phys}/H$, where $k_{\rm phys} = k/a$ is the physical wavenumber:

- *Sub-horizon* ($|k\tau| \gg 1$, i.e. $k_{\rm phys} \gg H$): $\Phi_k \sim e^{-ik\tau}$, oscillatory, behaving like a flat-space mode;
- *Horizon-crossing* ($|k\tau| \sim 1$): the transition regime;
- *Super-horizon* ($|k\tau| \ll 1$, i.e. $k_{\rm phys} \ll H$): $\Phi_k \to H/\sqrt{2k^3}$, a *frozen* amplitude that no longer oscillates.

The freezing of super-horizon modes is the decisive structural input. A mode that does not oscillate does not generate time-dependent decoherence in the way an oscillating mode does.

The Hubble scale thus separates modes that decohere from modes that do not—the infrared regulator the flat-space calculation could not provide.

2.3 The decoherence functional and its form factor

The coherent-state argument that underlies Eq. (7) is purely algebraic: a classical source coupled linearly to a free field displaces the vacuum into a coherent state, on any background [2]. What changes on de Sitter is not the structure of Eq. (7) but the *spectrum* of modes the displacement populates, since the mode functions are now Eq. (11) rather than plane waves.

For a source switched on at the start of the de Sitter epoch, the amplitude difference evolves as $\alpha_L(k,\tau) - \alpha_R(k,\tau) = \Delta\alpha^{\rm eq}(k)\, T(k,\tau,\tau_0)$, where $\Delta\alpha^{\rm eq}$ is the equilibrium (Coulomb-like) amplitude difference and T is a transient factor that interpolates from $T = 0$ at the switch-on time to $T \to 1$ once the field is established. Solving the sourced version of Eq. (10) exactly with the retarded Green's function built from the Bunch–Davies homogeneous solutions yields the transient factor in closed form,

$$T(k,\tau,\tau_0) = 1 - \frac{1}{a(\tau)}\, e^{ik(\tau_0-\tau)}, \tag{12}$$

the exponential-integral pieces cancelling identically. The decoherence-relevant part of $|T|^2$—obtained after subtracting the k-independent piece, which contributes only to the overall phase of the coherent-state overlap—is

$$|T|^2 - \left(1 - \tfrac{1}{a}\right)^2 = \frac{2}{a}\,(1 - \cos k\,\Delta\tau), \qquad \Delta\tau = \tau - \tau_0. \tag{13}$$

For sub-horizon modes Eq. (13) reduces, as $H \to 0$, exactly to the Minkowski factor $2(1-\cos\omega_k t)$—a derived consistency check. For super-horizon modes it is $O\big((k\,\Delta\tau)^2\big)$ and vanishes as $k \to 0$, which is the infrared suppression made explicit.

Inserting Eqs. (11) and (12) into the squared-norm exponent (7), and using physical (proper) separation $d = a(\tau)\,r$ in the angular integral, gives the de Sitter decoherence exponent as a mode integral

$$\Gamma_{\rm dS}(t) = \frac{16\,G^2\,\Delta m^2}{\pi\,\hbar\,c}\int_0^\infty \frac{dk}{k^3}\left(1 - \frac{\sin kd}{kd}\right)\left[\,|T(k,\tau)|^2 - (1-\tfrac{1}{a})^2\,\right]. \tag{14}$$

The angular factor $1-\sin(kd)/(kd)$, which is common to the Minkowski and de Sitter calculations, vanishes as $(kd)^2/6$ at small k and tends to unity for $kd \gg 1$. On Minkowski the integral (14) extends to $k = 0$; on de Sitter the renormalized transient factor cuts off the contribution of modes below the comoving horizon scale $k_H = aH/c = 1/R_{\rm H}^{\rm comoving}$, because those modes are frozen.

In the late-time, weak-coupling regime $d/c \ll t$, the sub-horizon part of Eq. (14) grows linearly in t, and the exponent takes the factorized form

$$\Gamma_{\rm dS}(t,d,H) = \frac{G\,\Delta m^2}{\hbar\,d}\, g\!\left(\frac{Hd}{c}\right) t. \tag{15}$$

The dimensionless form factor g is the ratio of the de Sitter gravitational energy scale to the Minkowski one, with the super-horizon modes removed by the horizon cutoff. Evaluating the cut-off mode integral in closed form (Appendix B) gives

$$\boxed{g(x) = 1 - \frac{2}{\pi}\,\mathrm{Si}(x)} \qquad \mathrm{Si}(x) = \int_0^x \frac{\sin u}{u}\,du. \tag{16}$$

The conversion of the saturating free-field mode integral into a rate that grows linearly in t—the step that takes Eq. (14) to Eq. (15)—is not automatic. It is the same constraint-driven step that underlies the G^1 scaling of Paper K: the influence functional alone gives a bounded, saturating exponent, and it is the Wheeler–DeWitt Hamiltonian constraint that converts it into linear-in-t growth. This point recurs in Sections 4.3 and 5; we flag here that the form factor g governs the *energy scale*, while the constraint governs the *rate*.

2.4 Limits: Minkowski recovery and the horizon

The form factor (16) has three regimes, and each carries physical content.

Minkowski limit, $Hd/c \to 0$. Here $\mathrm{Si}(0) = 0$, so $g(0) = 1$ exactly. The horizon cutoff $k_H = 1/R_{\mathrm{H}} \to 0$, the Bunch–Davies modes (11) reduce to plane waves, and the de Sitter calculation collapses to the Paper A/K result $\Gamma = G\,\Delta m^2/(\hbar d)$. This is the consistency check: had it failed, the de Sitter calculation would be wrong. The small-x onset is linear, $g(x) = 1 - (2/\pi)x + O(x^3)$, so the leading correction to the laboratory formula is suppressed by Hd/c—utterly negligible for any laboratory d, where $Hd/c \sim 10^{-26}$.

Horizon limit, $d \to R_{\mathrm{H}}$. Here $x = Hd/c = 1$, and

$$g(1) = 1 - \frac{2}{\pi}\,\mathrm{Si}(1) = 0.3977\ldots \tag{17}$$

The rate at the horizon scale is therefore

$$\Gamma_{\mathrm{dS}}^{\mathrm{rate}}(R_{\mathrm{H}}) = \frac{G\,\Delta m^2}{\hbar\,R_{\mathrm{H}}}\,g(1) = g(1)\,\frac{G\,\Delta m^2}{\hbar\,c}\,H, \tag{18}$$

finite—in sharp contrast with the divergent Minkowski extrapolation. Comparing to the causal bound,

$$\frac{\Gamma_{\mathrm{dS}}^{\mathrm{rate}}(R_{\mathrm{H}})}{H} = g(1)\,\frac{G\,\Delta m^2}{\hbar\,c} = g(1)\left(\frac{\Delta m}{m_P}\right)^2, \tag{19}$$

with $m_P = \sqrt{\hbar c/G}$ the Planck mass. For any sub-Planckian mass difference the horizon-scale rate is comfortably sub-causal.

Super-horizon limit, $Hd/c \gg 1$. Here $\mathrm{Si}(x) \to \pi/2$ and $g(x) \to 0$. Mass configurations separated by more than the Hubble radius produce no linear-in-t decoherence at all: their which-path information lies in frozen, super-horizon modes. This is the de Sitter causal horizon expressing itself directly in the decoherence functional—configurations outside each other's causal patch cannot dynamically decohere one another.

2.5 What the de Sitter recomputation does, and does not, settle

The de Sitter recomputation removes the *structural* pathology of W2. The catastrophic 10^{103} Hz of Eq. (2) was an artifact of using the Minkowski propagator—a propagator with no infrared scale—at a separation where the only physically correct propagator, the de Sitter one, has a horizon. With the horizon cutoff in place the rate is finite, proportional to H, and bounded by Eq. (19).

It does *not*, by itself, settle the cosmic rate. Equation (18) is finite but still contains the free input Δm. Inserting $\Delta m = M_U$ returns

$$\Gamma_{\rm dS}(M_U, R_{\rm H}) \sim g(1)\,\frac{G\,M_U^2}{\hbar c}\,H \sim 10^{103}\ \text{Hz}, \tag{20}$$

the same catastrophic number—because $(M_U/m_P)^2 \sim 10^{122}$ overwhelms the form factor. The de Sitter background fixes the propagator; it does not fix the input. Resolving W2 completely requires identifying what Δm *should* be for the cosmic question. That is the subject of Section 3.

3 The cosmic mass scale

The de Sitter form factor gives a rate per unit Δm^2 but leaves Δm open. This section identifies the physically correct Δm for cosmic decoherence and assembles the total rate. The conclusion is that Δm is not a free parameter: the structure of branches in quantum cosmology fixes it.

3.1 Why M_U is the wrong input

The cosmic wavefunction Ψ obeys the Wheeler–DeWitt constraint $\hat{H}_{\rm total}\,|\Psi\rangle = 0$. There is no external time; physical time is relational, recovered through the decoherent-histories construction [16, 17]. In that construction a coarse-graining of configuration space—here the spatial mass distribution smeared over horizon-scale cells—defines a set of histories, and the histories acquire probabilities, and a meaningful "rate," only once they decohere.

A *branch* of the cosmic wavefunction is then an equivalence class of Wheeler–DeWitt configurations that the decoherent-histories coarse-graining cannot tell apart. This is a definition, and it is the operationally meaningful one: it is what "physically distinct branch" *means*. The branch *granularity* Δm_* is the resolution of that coarse-graining—the smallest mass-distribution difference that the decoherence functional resolves into distinct, mutually decohered histories.

A configuration labelled "all of M_U displaced by $R_{\rm H}$" is not adjacent to the fiducial branch in history space. It is separated from it by an enormous number of elementary branch steps. Inserting $\Delta m = M_U$ into Eq. (18) computes the decoherence rate between two maximally distant branches—which is not the cosmic decoherence rate any more than coherently displacing every molecule of a gas by one metre defines that gas's decoherence rate. The physically meaningful cosmic rate is built from the *elementary* branch step Δm_*, summed over the independent branch-distinguishing degrees of freedom. We now determine both factors.

3.2 The branch granularity

The branch-distinguishing degrees of freedom are the graviton modes already identified by the de Sitter calculation: the Bunch–Davies modes crossing the horizon, each carrying one independent quantum of which-branch information. We determine the granularity Δm_* per such mode. Several independent lines of physical argument bear on it, and they converge.

Route A: thermal distinguishability floor. A static observer in de Sitter sees the Bunch–Davies vacuum as a thermal state at the Gibbons–Hawking temperature [9],

$$T_{\rm dS} = \frac{\hbar H}{2\pi k_B}. \tag{21}$$

Two cosmic branches whose mass-energy in a given horizon mode differs by less than the thermal energy fluctuation of that mode are not operationally distinguishable: the difference is buried inside the thermal ensemble. The distinguishability floor per mode is the thermal mass-energy spread, $\Delta E_{\rm floor} \sim k_B T_{\rm dS}$, giving $\Delta m_* \sim k_B T_{\rm dS}/c^2$. (We note an honest caveat: the single-mode fluctuation formula gives $k_B T_{\rm dS}$ exactly only in the long-wavelength equipartition limit; the strict horizon mode $x = \hbar\omega/k_B T_{\rm dS} = 2\pi$ gives $0.26\, k_B T_{\rm dS}$. The natural canonical value, and the one Routes B and C deliver sharply, is $k_B T_{\rm dS}$; an $O(1)$ factor in the range ∼0.3–1 is defensible and propagates to an $O(1)$ uncertainty in $\Gamma_{\rm total}$, not to its scaling.) One subtlety deserves note: de Sitter *spacetime* has negative heat capacity, which would invalidate a canonical fluctuation argument. But the fluctuation argument here is applied to the bulk *field modes*, which are ordinary harmonic oscillators in a genuine positive-temperature Gibbs state; the negative heat capacity is a property of the gravitating background, not of the matter sector, exactly as for Hawking radiation from a black hole.

Route B: the de Sitter horizon first law. The de Sitter horizon obeys a first law $dE = T_{\rm dS}\, dS_{\rm horizon}$ [9, 10, 11]. One quantum of horizon entropy—the smallest distinguishable change, $dS_{\rm horizon} = k_B$—corresponds to an energy $dE = k_B T_{\rm dS}$, hence a mass quantum $\Delta m_* = k_B T_{\rm dS}/c^2$. This is a genuinely distinct argument: the holographic principle native to this framework (Paper B) says the horizon stores $S_{\rm dS}$ which-branch registers, and the first law converts one register flip—one branch distinction—into one quantum of mass-energy $k_B T_{\rm dS}/c^2$.

Route C: the predictability sieve. Zurek's predictability sieve [20, 21] selects, among candidate branch spacings, those that the environment can hold a stable record of. Branches finer than one thermal quantum of the horizon environment cannot acquire a stable record—any incipient record is thermally erased on the de Sitter mixing time $\sim 1/H$. Branches at the thermal quantum are imprinted, frozen as the recording mode redshifts to super-horizon, and thereafter permanent. The einselected branch cell has width $\Delta m_* = k_B T_{\rm dS}/c^2$.

Convergence, and the irreducible premise. Routes A, B and C land on the same value,

$$\Delta m_{\rm dS} = \frac{k_B T_{\rm dS}}{c^2} = \frac{\hbar H}{2\pi c^2}. \tag{22}$$

Three arguments from genuinely different physics—bulk QFT, horizon thermodynamics, decoherence theory—reaching the same $O(1)$ coefficient, including the 2π, is informative. But the convergence must be read honestly. All three funnel through one shared premise: that the de Sitter environment is a Gibbs state at $T_{\rm dS}$. It is one thermodynamic insight reached three ways, not three independent results.

A fourth route, constructed without thermality, sharpens the picture. It reaches Eq. (22) from the Wheeler–DeWitt mini-superspace branch structure, the canonical clock–energy pairing $[\tau, E_{\text{static}}] = i\hbar$, and the energy–time uncertainty floor $\Delta E_* = \hbar/\beta_{\text{dS}}$—using no Gibbs state, no occupation number, no temperature. What this fourth route cannot avoid, and what Routes A–C also rely on, is the strictly weaker, *geometric* fact that Euclidean de Sitter has period

$$\beta_{\text{dS}} = \frac{2\pi}{H}. \tag{23}$$

Euclidean de Sitter is the round four-sphere of radius c/H; the period (23) is a property of that manifold. Thermality is *derived* from Eq. (23) via the KMS theorem, not the reverse. The honest statement of the convergence is therefore: the four routes do not share the premise of thermality—three assume it, one does not—but all four share the premise (23), a fact about the classical de Sitter background geometry. The cosmic-decoherence sector rests on that geometric premise, not on the hypothesis that the universe is in a hot thermal state. We regard the granularity (22) as physically well-motivated and reduced to the single geometric premise (23); we do not claim it is derived with the rigor of a theorem.

3.3 The holographic mode count

The cosmic wavefunction does not have infinitely many branch-labelling degrees of freedom. The holographic principle, native to this framework through Paper B, caps the information content of the de Sitter causal diamond at the Gibbons–Hawking entropy,

$$S_{\text{dS}} = \frac{A_H}{4\,\ell_P^2} = \frac{\pi\,c^5}{G\,\hbar\,H^2}, \tag{24}$$

with $\ell_P^2 = G\hbar/c^3$ the Planck area. This is the maximum number of independent which-branch registers the horizon can hold; finer modes would have to be stored on sub-Planckian horizon area, which the holographic bound forbids. The sum over branch-distinguishing modes therefore runs over exactly $N_{\text{modes}} = S_{\text{dS}}$ modes. Equivalently: the de Sitter infrared cutoff $k_H = 1/R_{\text{H}}$ removes super-horizon modes, the Planck scale cuts the ultraviolet end, and the number of independent modes in between is the horizon area in Planck units. The mode-counting statement and the holographic statement are the same statement. For $H = H_0$, Eq. (24) gives $S_{\text{dS}} \sim 2.5 \times 10^{122}$, consistent with standard estimates of the entropy of the observable universe [13].

3.4 Assembling the total rate

The per-mode rate follows from Eq. (18) with $\Delta m = \Delta m_{\text{dS}}$—each thermal-mass branch quantum is itself a horizon-scale object, $d = R_{\text{H}}$, and so carries the horizon form factor $g(1)$:

$$\Gamma_{\text{per mode}} = g(1)\,\frac{G\,\Delta m_{\text{dS}}^2\,H}{\hbar\,c} = g(1)\,\frac{G\,\hbar\,H^3}{4\pi^2\,c^5} = g(1)\,\frac{t_P^2\,H^3}{4\pi^2}, \tag{25}$$

using $t_P^2 = G\hbar/c^5$. Numerically, for $H = H_0$, this is $\sim 4 \times 10^{-142}$ Hz per mode—each individual mode decoheres on a timescale enormously longer than the Hubble time, as it must, since a single mode carries only one thermal quantum's worth of mass difference.

The total cosmic decoherence rate is the sum of the per-mode rates. In the free-field (Gaussian) approximation the Bunch–Davies state factorizes mode by mode and the linearized graviton coupling is mode-diagonal, so the total decoherence exponent is the sum of per-mode exponents,

$$\Gamma_{\text{total}}(t) = \sum_{i=1}^{S_{\text{dS}}} \Gamma_i(t). \tag{26}$$

We stress that this is the sum of S_{dS} *independent* two-branch decoherence processes, each with the *same* elementary granularity Δm_{dS}—it is not a partition of M_U into S_{dS} pieces. The Δm^2 non-linearity in Eq. (1) acts within one mode; it does not couple distinct modes. (The validity of Eq. (26) beyond strict Gaussian order, and the smallness of the mode-mode coupling correction, are taken up in Section 4.1.)

Combining Eqs. (25), (24) and (26),

$$\Gamma_{\text{total}} = S_{\text{dS}} \cdot \Gamma_{\text{per mode}} = \frac{\pi c^5}{G \hbar H^2} \cdot g(1) \frac{G \hbar H^3}{4\pi^2 c^5} = g(1) \frac{H}{4\pi}. \tag{27}$$

The cancellation is exact: G, $\hbar$ and c all drop out, and the result depends only on H and on the form factor $g(1)$. The bare algebraic identity $S_{\text{dS}} \cdot t_P^2 H^3/(4\pi^2) = H/(4\pi)$ is exact and $g(1)$-independent—it is the $g(1) = 1$ idealization, the relation between the holographic mode count and the per-mode-rate *formula*. The physical rate carries the form factor:

$$\Gamma_{\text{total}} = g(1) \frac{H}{4\pi} = \left[1 - \tfrac{2}{\pi} \operatorname{Si}(1)\right] \frac{H}{4\pi} \approx \frac{H}{10\pi} \approx 7 \times 10^{-20} \text{ Hz} \quad (H = H_0). \tag{28}$$

Equivalently $\Gamma_{\text{total}} t_H = g(1)/(4\pi) \approx 0.032$: the observable universe undergoes a fraction $\approx 1/30$ of a cosmic decoherence event per Hubble time.

Several internal checks support Eq. (28). *Dimensions*: $[\Delta m_{\text{dS}}] = [\hbar H/c^2] = \text{kg}$, $[\Gamma_{\text{per mode}}] = [t_P^2 H^3] = \text{s}^{-1}$, $[\Gamma_{\text{total}}] = [H] = \text{s}^{-1}$, with $g(1)$ dimensionless. *Causality*: $\Gamma_{\text{total}} = g(1) H/(4\pi) < H = c/R_{\text{H}}$, sub-causal by the factor $4\pi/g(1) \approx 10\pi$—with $g(1) < 1$, even further below the bound than the $g(1) = 1$ idealization. *Minkowski limit*: as $H \to 0$, $\Delta m_{\text{dS}} \to 0$ and $S_{\text{dS}} \to \infty$, but the product $\Gamma_{\text{total}} = g(1)H/(4\pi) \to 0$—no cosmic decoherence in flat space, as required, since flat space has no horizon and no branch granularity. *Laboratory consistency*: for a laboratory superposition the relevant Δm is the actual mass difference, vastly larger than $\Delta m_{\text{dS}} \sim 10^{-69}$ kg; the cosmic granularity is a floor, irrelevant when an experiment resolves branches far coarser than it, so the laboratory formula (1) with $g(0) = 1$ is recovered without conflict.

These relations have been checked numerically; an independent verification confirms the $g(1)$-independent algebraic identity to machine precision and the value $\Gamma_{\text{total}} \approx 7 \times 10^{-20}$ Hz at $H = H_0$ [5].

3.5 W2 resolved

W2 was the statement that the naive cosmic rate $\sim 10^{103}$ Hz violates causality by 121 orders of magnitude. Its resolution has the two parts this section and the previous one supply:

- *Structural* (Section 2): the Minkowski formula is the wrong propagator at $d \sim R_{\text{H}}$. The de Sitter propagator with its horizon cutoff $k_H = 1/R_{\text{H}}$ gives a finite rate, proportional to Δm^2 and to H.

- *Granularity* (this section): the cosmic Δm is not M_U. It is the de Sitter thermal mass $\Delta m_{\rm dS} = \hbar H/(2\pi c^2)$ per horizon mode. The $S_{\rm dS}$ horizon modes decohere independently and their rates add, giving $\Gamma_{\rm total} = g(1)\, H/(4\pi)$.

The 10^{103} Hz number came from inserting $\Delta m = M_U$, which the quantum-cosmology analysis identifies as the wrong physical input—a maximally distant branch jump, not a branch granularity. With the propagator corrected and the granularity identified, the cosmic decoherence rate is finite, proportional to H, and manifestly sub-causal. The internal contradiction is removed.

4 Internal consistency

The result of Section 3 was obtained in a controlled approximation: linearized gravity, the Bunch–Davies Gaussian vacuum, a scalar proxy for the graviton, and a particular extraction of the rate from the saturating mode integral. A result that resolves an internal contradiction must not introduce new ones, so each of these approximations was examined for whether it could affect the cosmic rate (28). This section reports the outcome of that examination. None of the four checks changes the rate; each leaves a stated residual, and we record those residuals here rather than in the discussion, so that the reader can see exactly what the result rests on.

4.1 Linear additivity of per-mode rates

Equation (26)—the statement that the $S_{\rm dS}$ horizon modes decohere into independent environments whose decoherence exponents add—is exact at strict Gaussian order. Mode–mode coupling enters only through graviton self-interaction, which is $O(G^2)$, parametrically subleading to the $O(G)$ leading coupling, exactly as for the G^1 scaling itself.

That subleading correction was computed rather than merely power-counted. An explicit leg-by-leg counting of powers of G in the connected two-vertex cut diagram finds the per-pair correction to be $O(\epsilon_{\rm dS})$, with $\epsilon_{\rm dS} = G\hbar H^2/(2\pi c^5) = 1/(2S_{\rm dS})$ the de Sitter loop-counting parameter—one power of $\epsilon_{\rm dS}$, not the $\epsilon_{\rm dS}^2$ a naive estimate would suggest, because $\epsilon_{\rm dS}$ is itself $O(G)$. One power is not, on its own, enough to defeat the combinatorial $S_{\rm dS}^2$ pair count: a coherent same-sign sum over all pairs would give a correction of order $S_{\rm dS}\, \epsilon_{\rm dS} = O(1)$, and linear additivity would fail.

What saves additivity is a sign. The per-pair matrix element carries a factor $\cos\theta_{ij}$, so the signed pair sum is a structure-factor sum—and a structure-factor sum is generically incoherent. Writing $\chi = \langle e^{i\psi}\rangle$ for the single-mode characteristic function of the combined configuration-plus-mode-label phase, the correction evaluates to

$$\frac{\delta\Gamma_{\rm total}}{\Gamma_{\rm total}} \sim \frac{c_v}{4}\, \max\left(|\chi|^2,\ 1/S_{\rm dS}\right), \tag{29}$$

with $c_v \approx 0.18$ a computed $O(1)$ coefficient. Linear additivity holds, $\delta\Gamma_{\rm total}/\Gamma_{\rm total} \ll 1$, if and only if $|\chi| \ll 1$. Computing $|\chi|$ directly from the de Sitter mode spectrum—the holographic horizon modes occupy a band in comoving wavenumber of width $\sim k_H\, S_{\rm dS}^{1/3}$, and the mode-label phase is the exact Bunch–Davies phase at the readout time—gives a smooth, closed-form characteristic function with $|\chi| \sim 3\, S_{\rm dS}^{-1/3} \sim 5 \times 10^{-41}$. Linear additivity therefore holds with roughly forty orders of margin, and the honest residual magnitude is $\delta\Gamma_{\rm total}/\Gamma_{\rm total} \sim (c_v/4)\, |\chi|^2 \sim 2 \times 10^{-82}$.

This result is conditional on one mild premise: that the decoherent-histories readout conformal time is not fine-tuned to lie below horizon crossing ($|\tau_*| \gtrsim 1/k_H$). We label this premise P-readout and carry it forward to Section 5; it is a one-sided, geometric quantum-cosmology premise, and a derivation of the readout time is the natural next step rather than an obstruction to the present result.

4.2 The spin-2 graviton

The derivation used a scalar proxy: each transverse-traceless graviton polarization treated as a massless minimally coupled scalar. Redoing the key computations with the full spin-2 graviton $h_{\mu\nu}^{\mathrm{TT}}$, its two polarizations, and the appropriate tensor harmonics on de Sitter [8] confirms the proxy. The transverse-traceless graviton on de Sitter obeys, per polarization, *exactly* the scalar-proxy mode equation (10) [7]. Consequently the form factor $g(x) = 1 - (2/\pi)\,\mathrm{Si}(x)$ and its horizon value $g(1) = 0.398$ carry a spin-2 rescaling factor of exactly 1.000. The only quantitative refinement is the per-pair coefficient of Section 4.1, rescaled by the polarization factor 8/15 from $c_v = 0.34$ (scalar proxy) to $c_v = 0.18$—a change in a safe direction that does not affect the power of ϵ_{dS}. The cosmic rate (28) is unchanged. The scalar proxy is not an approximation here; for the mode structure that governs the form factor it is exact.

4.3 Gauge-invariant rate extraction

A decoherence "rate" computed on a cosmological background risks depending on the choice of time slicing. As initially formulated, the proper-time and conformal-time observers appeared to give different $O(1)$ coefficients, a spread $g(1) \in [0.21, 0.43]$.

This is resolved by extracting the rate from an intrinsic, foliation-free clock. The de Sitter static-patch crossed-product algebra of observables [22, 23] is of Type II_1, with a finite, observer-independent trace; its modular flow is the static-patch Killing flow. Reading the decoherence rate off the modular-flow derivative of the modular relative entropy makes no reference to any cosmological time slicing. The gauge-invariant rate so obtained is

$$\Gamma_{\mathrm{inv}} = g(1)\,\frac{H}{4\pi}, \qquad g(1) = 1 - \frac{2}{\pi}\,\mathrm{Si}(1) = 0.398, \tag{30}$$

the energy-scale value of the form factor, pinned exactly. The earlier spread $[0.21, 0.43]$ is identified as an artifact of frame-dependent phase-averaging, which the modular extraction bypasses, and is eliminated.

One element of this extraction is not a gauge ambiguity but an irreducible input: the rate is per unit comoving proper time with the surface-gravity normalization $\kappa = H$—which is again the de Sitter Euclidean period $2\pi/H$, the same geometric premise (23) that the granularity rests on. de Sitter homogeneity makes all comoving observers equivalent, so this is not a residual ambiguity; it is a framework input, and it is the *same* input already on the table.

4.4 All-orders secular control

The graviton on de Sitter is known to have an infrared sector with secular growth: graviton self-interaction produces loop corrections growing as $\ln a$ at each order in perturbation theory [24, 25]. One must check that this does not invalidate the linearized calculation on the timescale of interest.

The first point is structural: the secular factor $\ln a$ multiplies *only* the $O(G^2)$ mode–mode coupling correction of Section 4.1—never the leading additive sum (27), which is built from the Gaussian Bunch–Davies state and the mode-diagonal linearized coupling and carries no self-interaction loop. The leading cosmic rate is untouched by secular growth at any order.

The second point concerns the timescale. The relevant question is the rate at which the cosmic wavefunction as a whole decoheres, and that is the *sum* over $S_{\rm dS}$ modes; the sum reaches $O(1)$ after only a few Hubble times, $t_{\rm dec} \sim 4\pi/H$, even though any single mode would take $\sim 10^{123}$ Hubble times in isolation. The per-mode time is a counterfactual that never enters, because all $S_{\rm dS}$ channels decohere in parallel. The secular factor must therefore be evaluated at $t \sim t_{\rm dec}$, where $\ln a = O(1)$.

This timescale argument can be promoted to an all-orders bound. The leading-log tower $L(x) = 1 + \sum_n c_{n,n}\, x^n$ has secular argument $x = \kappa^2 H^2 \,\ln a$ with $\kappa^2 H^2 \sim \epsilon_{\rm dS}$; its $O(1)$ coefficients give a geometric bound $|L(x) - 1| \le C\,|x|/(1 - |x|)$. On the cosmic-decoherence timescale the argument is $x_{\rm dec} \sim 5 \times 10^{-122}$, so the all-orders secular correction to the $O(G^2)$ additivity term is $\le 10^{-203}$; the crossover at which it could reach $O(1)$ lies ~ 120 orders of magnitude beyond $t_{\rm dec}$. This bound uses only the established leading-log structure and the gauge-invariant readout time of Section 4.3—no resummation is needed. A Starobinsky/dynamical-renormalization-group resummation strengthens it further, indicating the leading-log tower sums to a bounded equilibrium-type function rather than a divergent series.

Two residuals survive: the full spin-2 all-orders graviton resummation, and the complete subleading-log resummation. Both are open problems of the quantum-field-theory-on-de-Sitter literature, not of QGD; and the timescale bound, which does not resum the series, shows that neither can affect the cosmic rate—on $t \sim t_{\rm dec}$ the secular argument is $x \sim 10^{-121}$ regardless of how the literature resolves them.

4.5 Status after the consistency checks

Table 1 collects the four checks. Each was passed; the cosmic rate $\Gamma_{\rm total} = g(1)\,H/(4\pi)$ is unchanged by all of them. What remains is a small, explicitly named set of premises—the de Sitter Euclidean period (23), the mild geometric premise P-readout, the holographic finiteness of the mode count, the decoherent-histories definition of relational time, and the framework-wide energy-to-rate identification—and it is to these that the discussion turns.

Table 1. *Outcome of the internal-consistency checks on the cosmic decoherence rate* $\Gamma_{\text{total}} = g(1)\,H/(4\pi)$.

Approximation	**Finding**	**Residual**
Linear additivity (4.1)	Mode–mode coupling is $O(G^2)$; the signed pair sum is incoherent; $\delta\Gamma_{\text{total}}/\Gamma_{\text{total}} \sim 2\times 10^{-82}$	Conditional on premise P-readout (the readout time is not fine-tuned below horizon crossing)
Scalar proxy (4.2)	TT graviton obeys exactly the scalar mode equation; $g(1)$ rescaling is 1.000	None; proxy is exact for the form factor
Rate gauge (4.3)	Type II_1 modular flow gives a foliation-free rate; $g(1) = 0.398$ pinned	The normalization $\kappa = H$, i.e. the period $2\pi/H$ — the same geometric premise
Secular growth (4.4)	Dresses only the $O(G^2)$ term; bounded $\leq 10^{-203}$ on t_{dec}	Full spin-2 / subleading-log resummation — literature problems, cannot affect the rate

5 Discussion

This section states plainly what the result is, what it rests on, and what would overturn it. It is written in the conviction that a resolution of an internal contradiction is worth only as much as the honesty with which its remaining premises are declared.

5.1 What is established

The structural claim is firm. The catastrophic cosmic rate of Eq. (2) is an artifact of using the Minkowski graviton propagator at separations $d \sim R_{\text{H}}$, where that propagator is not the correct Green's function. This is not a matter of interpretation: the Minkowski propagator has no infrared scale, the de Sitter propagator has a horizon, and at $d \sim R_{\text{H}}$ the difference is the whole effect. Recomputing the coherent-state decoherence functional on a de Sitter background produces the form factor $g(x) = 1 - (2/\pi)\,\text{Si}(x)$, which is finite for all x, recovers the Paper A/K Minkowski result as $x \to 0$, and vanishes for super-horizon separations. The divergence of W2 is gone, and gone for a stated structural reason.

The quantitative claim is more qualified, and we are explicit about the qualification. Given the de Sitter form factor, the holographic mode count (24), the branch granularity (22), and linear additivity, the cosmic rate is $\Gamma_{\text{total}} = g(1)\,H/(4\pi) \approx 7\times 10^{-20}$ Hz by arithmetic. The form factor $g(1) = 0.398$ is computed in closed form and pinned gauge-invariantly. The granularity Δm_{dS} is reduced—through four converging routes—to the single geometric premise that Euclidean de Sitter has period $2\pi/H$. Linear additivity is established conditional on the mild premise P-readout. The internal-consistency program of Section 4 closes the scalar-proxy and secular-growth approximations outright. What remains is a residual $O(1)$ uncertainty in the granularity coefficient (the Route-A mode-weighting ambiguity, a factor in $\sim$0.3–1) and the named premises below.

We are careful about the verbs. The structural cure of W2 *is shown.* The de Sitter form factor *is derived* from the coherent-state functional. The granularity Δm_{dS} *is reduced to* a geometric premise—it is well-motivated and convergent, not proven with the rigor of a theorem.

The total rate $\Gamma_{\text{total}} = g(1)\, H/(4\pi)$ *follows* from those inputs. We do not claim to have proven the cosmic rate; we claim to have removed an internal inconsistency and replaced a 121-order causality violation with a finite, sub-causal, premise-labelled result.

5.2 The residuals, stated openly

Three classes of residual survive. We list them because a reader is entitled to know exactly where the result is soft.

The P-readout premise. Linear additivity (Section 4.1) holds because the single-mode characteristic function $|\chi|$ is small, and $|\chi| \sim 5 \times 10^{-41}$ was computed from the de Sitter mode spectrum conditional on the readout conformal time not being fine-tuned below horizon crossing. This is a mild, one-sided, geometric premise about the quantum-cosmological coarse-graining. It is not closed here; a derivation of the readout time—for which the Type II_1 crossed-product framework of Section 4.3 is the natural setting—is the cleanest piece of unfinished business specific to this paper.

The energy-to-rate identification. The whole QGD decoherence program, beginning with Paper A, rests on the identification of a gravitational energy scale with a decoherence rate: the influence functional supplies a saturating energy scale, and a Hamiltonian constraint converts it into a linearly growing rate, $\Gamma = E_G/\hbar$ rather than the bare modular gap frequency or its square. The de Sitter calculation inherits this identification; it does not establish it. This is the central G^1 assumption of the framework, it is not specific to de Sitter or to cosmic scale, and it is the deepest open question on which the present result—like Paper A and Paper K—depends. We name it here so that it is not mistaken for something this paper closes.

Literature-level open problems. The all-orders secular control of Section 4.4 leaves two residuals—the full spin-2 all-orders graviton resummation and the complete subleading-log resummation—that are open problems of quantum field theory on de Sitter space, not of QGD. The timescale bound shows that neither can affect the cosmic rate, because on the relevant timescale the secular argument is $\sim 10^{-121}$ regardless of how those problems are resolved. They are residuals of the surrounding literature that the present result is insulated from, not residuals of the present result.

A reader weighing the result should also keep in view the largest caveat of all, which is not a residual of this calculation but its setting: QGD is a framework without direct experimental confirmation. This paper resolves a contradiction *internal* to that framework. It strengthens QGD by removing a self-inconsistency; it does not, and cannot, establish QGD. The cosmic decoherence rate $\Gamma_{\text{total}} = g(1)\, H/(4\pi)$ is a prediction *of* QGD, conditional on QGD.

5.3 Relation to prior work

Decoherence in quantum cosmology has a long history. The recognition that the universe's wavefunction acquires classical behaviour through the decoherence of its coarse-grained histories goes back to Zeh [14], Kiefer [15], and Halliwell [16, 18], and the decoherent-histories formalism we use to define a cosmic branch is that of Gell-Mann and Hartle [17]. Barvinsky, Kamenshchik and Kiefer [19] computed decoherence of the cosmological wavefunction by its field environment. The present paper does not reopen the foundations of that program; it adopts the decoherent-histories definition of a branch and asks a sharper, quantitative question—the branch *granularity*—that the QGD decoherence formula makes well posed.

The granularity itself draws on the Gibbons–Hawking thermodynamics of the de Sitter horizon [9], on the thermodynamic reading of the Einstein equation [10, 11], and on Zurek's einselection [20, 21]; the holographic mode count is the Gibbons–Hawking entropy used as a bound on which-branch information [12]. The gauge-invariant rate extraction uses the de Sitter algebra of observables of Chandrasekaran, Longo, Penington and Witten [22] and the crossed-product construction of Witten [23]. The graviton infrared sector and its secular growth are the subject of the Tsamis–Woodard program [24, 25, 26].

What is new here is not any one of those ingredients but their assembly into the resolution of a specific, quantitative contradiction. The contribution is to show that the QGD laboratory decoherence formula, taken to cosmic scale, produces a causality paradox; that the paradox is a propagator artifact; and that the corrected calculation, fed the branch granularity that quantum cosmology supplies, returns a finite sub-causal rate $\Gamma_{\text{total}} = g(1)\, H/(4\pi)$.

5.4 Falsifiability

A resolution of an internal inconsistency should still expose itself to refutation. There are three distinct ways the picture of this paper could be shown wrong.

By a calculation. The result makes a sharp internal claim: the cosmic decoherence rate of QGD is $\Gamma_{\text{total}} = g(1)\, H/(4\pi)$, with $g(1) = 1 - (2/\pi)\,\mathrm{Si}(1)$ a closed-form number. The chain leading to it—the de Sitter form factor, the per-mode rate (25), the holographic mode count, linear additivity—is explicit and reproducible. A computation showing that the de Sitter coherent-state functional does *not* yield $g(x) = 1 - (2/\pi)\,\mathrm{Si}(x)$, or that the signed mode–mode pair sum is in fact coherent (so that additivity fails and $\delta\Gamma_{\text{total}}/\Gamma_{\text{total}} = O(1)$), or that the branch granularity is parametrically different from $\hbar H/(2\pi c^2)$, would overturn the result. The most exposed link is the P-readout premise: were the decoherent-histories readout time forced below horizon crossing, $|\chi|$ would not be small and linear additivity would break. This is a concrete target for the next calculation.

By the laboratory. The cosmic result and the laboratory prediction of Paper A are not independent: the de Sitter form factor reduces to the laboratory formula (1) as $Hd/c \to 0$, with $g(0) = 1$ exact. The laboratory G^1 rate is testable, with current technology striving toward the required regime; observation of G^2 scaling rather than G^1 would falsify the constrained-influence-functional mechanism of Paper K. Because the cosmic calculation uses the *same* mechanism—the same constraint-extracted rate, the same coherent-state overlap, only on a de Sitter rather than a Minkowski background—a laboratory falsification of G^1 would also remove the basis of the cosmic result. The cosmic and laboratory predictions stand or fall together on the mechanism, even though the cosmic rate is far too slow to observe directly.

By cosmology. The predicted rate $\Gamma_{\text{total}} = g(1)\, H/(4\pi) \approx H/(10\pi)$ is too slow to act as a direct observational signal: a fraction $\sim 1/30$ of a decoherence event per Hubble time is not something a measurement resolves. Its content is instead a sharp internal statement—the universe decoheres at a definite fraction of the Hubble rate, with the coefficient fixed by $g(1)$ and the geometry. The result also makes a structural prediction that is in principle distinguishing: cosmic gravitational decoherence is set by the de Sitter scale H and not by the total mass M_U; configurations separated by more than a Hubble radius do not dynamically decohere one another. A demonstration that cosmic-scale superpositions decohere at a rate controlled by M_U rather than H, or across super-horizon separations, would contradict the picture. We do not overstate this: there is at present no experiment that probes cosmic-scale gravitational decoherence, and we claim none.

The honest summary is that the result's primary exposure is to calculation. It is a falsifiable claim about what QGD predicts, with an explicit derivation chain and one clearly marked soft link; it shares the laboratory program's exposure through the common G^1 mechanism; and it removes a contradiction that, left standing, would have been reason to doubt the framework.

6 Conclusion

The gravitational decoherence formula $\Gamma = G\,\Delta m^2/(\hbar d)$ of Quantum-Geometric Duality, applied without examination at cosmic scale, returns a rate $\sim 10^{103}$ Hz—121 orders of magnitude above the causal bound c/R_{H}. This was a genuine internal contradiction of the framework. This paper has resolved it.

The resolution is in two parts. The divergence is a propagator artifact: the formula is built on the Minkowski graviton two-point function, which has no infrared scale, and at separations approaching the Hubble radius that is the wrong Green's function. Recomputed on a de Sitter background, the coherent-state decoherence functional acquires a closed-form form factor $g(x) = 1 - (2/\pi)\,\mathrm{Si}(x)$, which recovers the laboratory result for $Hd/c \ll 1$, is finite at the horizon, and vanishes for super-horizon separations—the de Sitter horizon supplying the infrared cutoff that flat space lacked. Separately, the cosmic decoherence rate is not governed by the total mass M_U but by the elementary branch granularity of the cosmic wavefunction, identified by several converging arguments with the de Sitter thermal mass $\Delta m_{\mathrm{dS}} = \hbar H/(2\pi c^2)$ per horizon mode. Summed over the S_{dS} holographic horizon modes, the cosmic rate is

$$\Gamma_{\mathrm{total}} = g(1)\,\frac{H}{4\pi} \approx \frac{H}{10\pi} \approx 7 \times 10^{-20}\ \mathrm{Hz}, \tag{31}$$

finite, proportional to H, and comfortably sub-causal.

An internal-consistency program established that this rate is unchanged by the remaining approximations: the scalar proxy is exact for the graviton mode structure that governs the form factor; the rate is gauge-invariant under the de Sitter modular flow; linear additivity of per-mode rates holds with forty orders of margin; and the known de Sitter graviton secular growth dresses only an $O(G^2)$ correction and is bounded far below any level that could matter. The result rests on a small, named set of premises—chiefly the geometric de Sitter Euclidean period $2\pi/H$, a mild premise on the quantum-cosmological readout time, and the framework-wide identification of a gravitational energy scale with a decoherence rate—and we have stated each one openly.

The result strengthens QGD by removing a self-inconsistency; it does not establish the framework, which awaits experimental test. Its sharpest exposure is to calculation: the derivation chain is explicit and reproducible, with one clearly marked soft link, the readout-time premise, whose derivation is the natural next step. Through the common G^1 mechanism it also shares the falsifiability of the laboratory program of Paper A. What the paper provides is a finite, premise-labelled, causally consistent answer to a question the framework had previously left as an open contradiction: the universe decoheres gravitationally, into its own horizon, at a definite fraction of the Hubble rate.

Appendices

A Conventions and notation

This appendix collects the conventions used throughout the paper.

We work in SI units. The fundamental constants are Newton's constant $G = 6.674 \times 10^{-11}\ \mathrm{m^3 kg^{-1} s^{-2}}$, the reduced Planck constant $\hbar = 1.055 \times 10^{-34}\ \mathrm{J \cdot s}$, and the speed of light $c = 2.998 \times 10^8$ m/s. From these we form the Planck length $\ell_P = \sqrt{G\hbar/c^3} = 1.616 \times 10^{-35}$ m, the Planck time $t_P = \ell_P/c$ so that $t_P^2 = G\hbar/c^5$, and the Planck mass $m_P = \sqrt{\hbar c/G} = 2.176 \times 10^{-8}$ kg.

The cosmological inputs are the Hubble rate H, taken at its present value $H_0 \approx 2.2 \times 10^{-18}\ \mathrm{s^{-1}}$ for numerical estimates; the Hubble radius $R_{\mathrm{H}} = c/H$; the Hubble time $t_H = 1/H$; and the horizon area $A_H = 4\pi R_{\mathrm{H}}^2$. The de Sitter (Gibbons–Hawking) temperature is $T_{\mathrm{dS}} = \hbar H/(2\pi k_B)$, and the de Sitter entropy is $S_{\mathrm{dS}} = A_H/(4\ell_P^2) = \pi c^5/(G\hbar H^2)$.

We adopt the metric signature $(-,+,+,+)$. The de Sitter background is written in flat slicing as $ds^2 = -c^2\,dt^2 + a(t)^2\,d\mathbf{x}^2$ with $a(t) = e^{Ht}$, or in conformal time τ (defined by $a\,d\tau = c\,dt$) as $ds^2 = a(\tau)^2(-c^2\,d\tau^2 + d\mathbf{x}^2)$ with $a(\tau) = -1/(H\tau)$ and $\tau \in (-\infty, 0)$. The linearized graviton is a perturbation $h_{\mu\nu}$ of this background; in transverse-traceless gauge each polarization is treated, in the scalar proxy, as a massless minimally coupled scalar Φ.

The primary quantities of the decoherence calculation are the branch mass difference Δm, the proper separation $d = a(\tau)\,|\mathbf{x}_L - \mathbf{x}_R|$, the comoving wavenumber k, the physical wavenumber $k_{\mathrm{phys}} = k/a$, and the comoving horizon wavenumber $k_H = aH/c$. The dimensionless de Sitter parameter is $x = Hd/c = d/R_{\mathrm{H}}$. The decoherence exponent $\Gamma(t)$ enters the off-diagonal density-matrix element as $\rho_{LR}(t) = \rho_{LR}(0)\,e^{-\Gamma(t)}$; the decoherence *rate* is $\Gamma(t)/t$ in the linear-growth regime.

The sine integral is $\mathrm{Si}(x) = \int_0^x (\sin u/u)\,du$, with $\mathrm{Si}(0) = 0$ and $\mathrm{Si}(\infty) = \pi/2$. The de Sitter form factor is $g(x) = 1 - (2/\pi)\,\mathrm{Si}(x)$, with $g(0) = 1$, $g(1) = 0.3977\ldots$, and $g(x \to \infty) = 0$.

B The de Sitter form factor

This appendix derives the closed form (16) of the de Sitter form factor $g(x)$ and records the limiting behaviour used in the main text.

B.1 The form factor as an energy-scale ratio

The free-field de Sitter decoherence functional grows logarithmically, not linearly, in time; the linear-in-t decoherence *rate* is set by the gravitational energy scale E_G through the constraint-extraction mechanism—the influence functional supplies the energy scale, the Hamiltonian constraint supplies the rate. The form factor $g(Hd/c)$ of the factorized law (15) is accordingly the ratio of the de Sitter gravitational energy scale at separation d to its Minkowski counterpart.

The d-dependent part of the gravitational interaction energy of the two mass configurations is built from the Coulomb propagator $1/k^2$; its dependence on the separation is carried by the radial integral

$$E_G^{\mathrm{Mink}} \propto \int_0^\infty dk\,\frac{\sin kd}{kd} = \frac{\pi}{2d}. \tag{32}$$

On de Sitter, modes with comoving wavenumber $k < k_H = 1/R_{\mathrm{H}}$ are super-horizon: their transient factor is frozen and they do not contribute to the time-dependent energy scale. The de

Sitter energy scale is the same integral carrying the horizon infrared cutoff,

$$E_G^{\mathrm{dS}}(x) \;\propto\; \int_{k_H}^{\infty} dk\,\frac{\sin kd}{kd} \;=\; \frac{1}{d}\Big(\frac{\pi}{2} - \mathrm{Si}(x)\Big), \qquad x = k_H d = \frac{Hd}{c}, \tag{33}$$

where the evaluation uses $\int_x^\infty u^{-1}\sin u\,du = \pi/2 - \mathrm{Si}(x)$.

B.2 The closed form

The form factor is the ratio of (33) to (32):

$$g(x) \;=\; \frac{E_G^{\mathrm{dS}}(x)}{E_G^{\mathrm{Mink}}} \;=\; \frac{\pi/2 - \mathrm{Si}(x)}{\pi/2} \;=\; 1 - \frac{2}{\pi}\,\mathrm{Si}(x). \tag{34}$$

No free constant enters: the Minkowski normalization $g(0) = 1$ is automatic, since as $H \to 0$ the cutoff $k_H \to 0$ and (33) reduces identically to (32).

B.3 Limits

From Eq. (34) and the standard values of the sine integral:

- $\mathrm{Si}(0) = 0$, so $g(0) = 1$. This is the Minkowski limit; the de Sitter calculation reproduces the Paper A/K result $\Gamma = G\,\Delta m^2/(\hbar d)$ exactly.
- Small-x expansion: $\mathrm{Si}(x) = x - x^3/18 + O(x^5)$, hence

  $$g(x) = 1 - \frac{2}{\pi}\,x + \frac{1}{9\pi}\,x^3 + O(x^5). \tag{35}$$

  The leading correction to the laboratory formula is linear in Hd/c and, for any laboratory separation where $Hd/c \sim 10^{-26}$, entirely negligible.
- At the horizon, $\mathrm{Si}(1) = 0.946083\ldots$, so

  $$g(1) = 1 - \frac{2}{\pi}\,(0.946083\ldots) = 0.39768\ldots, \tag{36}$$

  the value used throughout as $g(1) \approx 0.398$.
- $\mathrm{Si}(\infty) = \pi/2$, so $g(x \to \infty) = 1 - (2/\pi)(\pi/2) = 0$. Super-horizon separations produce no linear-in-t decoherence.

The function g decreases monotonically from 1 to 0 as x runs from 0 to ∞, since $g'(x) = -(2/\pi)\sin(x)/x$ is negative for $0 < x < \pi$ and the subsequent oscillations of $\sin(x)/x$ are too small to reverse the trend; g is positive throughout because $\mathrm{Si}(x) < \pi/2$ for all finite x.

C The holographic mode count and the cancellation

This appendix records the algebra behind the cancellation in Eq. (27) and the numerical estimates quoted in the main text.

C.1 The exact cancellation

The per-mode rate (25) is, with $t_P^2 = G\hbar/c^5$,

$$\Gamma_{\text{per mode}} = g(1)\,\frac{G\,\Delta m_{\text{dS}}^2\,H}{\hbar\,c}, \qquad \Delta m_{\text{dS}} = \frac{\hbar H}{2\pi c^2}. \tag{37}$$

Substituting the granularity,

$$\Gamma_{\text{per mode}} = g(1)\,\frac{G\,H}{\hbar\,c}\cdot\frac{\hbar^2 H^2}{4\pi^2 c^4} = g(1)\,\frac{G\,\hbar\,H^3}{4\pi^2\,c^5} = g(1)\,\frac{t_P^2\,H^3}{4\pi^2}. \tag{38}$$

The holographic mode count (24) is

$$S_{\text{dS}} = \frac{A_H}{4\,\ell_P^2} = \frac{4\pi(c/H)^2}{4\,(G\hbar/c^3)} = \frac{\pi\,c^5}{G\,\hbar\,H^2}. \tag{39}$$

The product is

$$\Gamma_{\text{total}} = S_{\text{dS}}\cdot\Gamma_{\text{per mode}} = \frac{\pi\,c^5}{G\,\hbar\,H^2}\cdot g(1)\,\frac{G\,\hbar\,H^3}{4\pi^2\,c^5} = g(1)\,\frac{H}{4\pi}. \tag{40}$$

Every factor of G, $\hbar$ and c cancels. The combination $S_{\text{dS}}\cdot t_P^2 H^3/(4\pi^2) = H/(4\pi)$ is an exact algebraic identity, independent of the form factor; it is the $g(1) = 1$ idealization. The physical rate carries $g(1)$, giving $\Gamma_{\text{total}} = g(1)\,H/(4\pi)$.

This cancellation is not an accident of bookkeeping. It expresses the structural fact that the holographic mode count and the per-mode rate are built from the same horizon: the number of which-branch registers and the rate at which each is written are both set by the de Sitter scale, so their product depends only on H.

C.2 Numerical values

At the present epoch, $H = H_0 \approx 2.20\times10^{-18}\ \text{s}^{-1}$:

$$T_{\text{dS}} = \frac{\hbar H_0}{2\pi k_B} \approx 2.7\times10^{-30}\ \text{K}, \tag{41}$$

$$\Delta m_{\text{dS}} = \frac{\hbar H_0}{2\pi c^2} \approx 4\times10^{-69}\ \text{kg}, \tag{42}$$

$$S_{\text{dS}} = \frac{\pi c^5}{G\hbar H_0^2} \approx 2.5\times10^{122}, \tag{43}$$

$$\Gamma_{\text{per mode}} = g(1)\,\frac{t_P^2 H_0^3}{4\pi^2} \approx 4\times10^{-142}\ \text{Hz}, \tag{44}$$

$$\Gamma_{\text{total}} = g(1)\,\frac{H_0}{4\pi} \approx 7\times10^{-20}\ \text{Hz}. \tag{45}$$

The $g(1) = 1$ idealization would give $H_0/(4\pi) \approx 1.75\times10^{-19}$ Hz; the physical rate is smaller by the factor $g(1) \approx 0.40$. The causal bound is $c/R_{\text{H}} = H_0 \approx 2.2\times10^{-18}$ Hz, so $\Gamma_{\text{total}}/H_0 = g(1)/(4\pi) \approx 0.032$: the cosmic rate is sub-causal by a factor of roughly 30. The mode count $S_{\text{dS}} \approx 2.5\times10^{122}$ agrees with standard estimates of the entropy of the observable universe [13].

These identities and values have been checked numerically to machine precision by an independent verification, which confirms the $g(1)$-independent algebraic identity $S_{\text{dS}}\cdot t_P^2 H^3/(4\pi^2) = H/(4\pi)$, the form-factor value $g(1) = 0.3977$, and the physical rate $\Gamma_{\text{total}} = g(1)\,H/(4\pi)$ [5].

References

[1] M. Sperzel, *Gravitational Decoherence via Entanglement-Geometry: Testing the Diósi-Penrose Hypothesis*, Quantum-Geometric Duality Series, Paper A (2026).

[2] M. Sperzel, *Gravitational Decoherence at $O(G)$: The Wheeler-DeWitt Constraint on the Feynman-Vernon Influence Functional*, Quantum-Geometric Duality Series, Paper K (2026).

[3] M. Sperzel, *Cosmic Decoherence and the Arrow of Time in Quantum-Geometric Duality*, Quantum-Geometric Duality Series, internal derivation (2026).

[4] M. Sperzel, *Vacuum Decoherence and the Cosmological Constant in Quantum-Geometric Duality*, Quantum-Geometric Duality Series, internal derivation (2026).

[5] M. Sperzel, *Cosmic gravitational decoherence: numerical verification suite*, Quantum-Geometric Duality Series, verification code `test_cosmic_decoherence.py` (2026).

[6] T. S. Bunch and P. C. W. Davies, *Quantum field theory in de Sitter space: Renormalization by point-splitting*, Proc. R. Soc. Lond. A **360**, 117 (1978). doi:10.1098/rspa.1978.0060

[7] L. H. Ford and L. Parker, *Quantized gravitational wave perturbations in Robertson-Walker universes*, Phys. Rev. D **16**, 1601 (1977). doi:10.1103/PhysRevD.16.1601

[8] A. Higuchi, *Symmetric tensor spherical harmonics on the N-sphere and their application to the de Sitter group $SO(N, 1)$*, J. Math. Phys. **28**, 1553 (1987). doi:10.1063/1.527513

[9] G. W. Gibbons and S. W. Hawking, *Cosmological event horizons, thermodynamics, and particle creation*, Phys. Rev. D **15**, 2738 (1977). doi:10.1103/PhysRevD.15.2738

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

[11] T. Padmanabhan, *Thermodynamical aspects of gravity: new insights*, Rep. Prog. Phys. **73**, 046901 (2010). doi:10.1088/0034-4885/73/4/046901

[12] R. Bousso, *The holographic principle*, Rev. Mod. Phys. **74**, 825 (2002). doi:10.1103/RevModPhys.74.825

[13] C. A. Egan and C. H. Lineweaver, *A larger estimate of the entropy of the universe*, Astrophys. J. **710**, 1825 (2010). doi:10.1088/0004-637X/710/2/1825

[14] H. D. Zeh, *Emergence of classical time from a universal wavefunction*, Phys. Lett. A **116**, 9 (1986). doi:10.1016/0375-9601(86)90346-4

[15] C. Kiefer, *Continuous measurement of mini-superspace variables by higher multipoles*, Class. Quantum Grav. **4**, 1369 (1987). doi:10.1088/0264-9381/4/5/031

[16] J. J. Halliwell, *Decoherence in quantum cosmology*, Phys. Rev. D **39**, 2912 (1989). doi:10.1103/PhysRevD.39.2912

[17] M. Gell-Mann and J. B. Hartle, *Classical equations for quantum systems*, Phys. Rev. D **47**, 3345 (1993). doi:10.1103/PhysRevD.47.3345

[18] J. J. Halliwell, *How the quantum universe became classical*, Contemp. Phys. **46**, 93 (2005). doi:10.1080/0010751052000297588

[19] A. O. Barvinsky, A. Yu. Kamenshchik, and C. Kiefer, *Effective action and decoherence by fields in quantum cosmology*, Nucl. Phys. B **552**, 420 (1999). doi:10.1016/S0550-3213(99)00208-4

[20] W. H. Zurek, *Preferred states, predictability, classicality and the environment-induced decoherence*, Prog. Theor. Phys. **89**, 281 (1993). doi:10.1143/ptp/89.2.281

[21] W. H. Zurek, *Decoherence, einselection, and the quantum origins of the classical*, Rev. Mod. Phys. **75**, 715 (2003). doi:10.1103/RevModPhys.75.715

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

[23] E. Witten, *Gravity and the crossed product*, J. High Energy Phys. **2022**(10), 008 (2022). doi:10.1007/JHEP10(2022)008

[24] N. C. Tsamis and R. P. Woodard, *Quantum gravity slows inflation*, Nucl. Phys. B **474**, 235 (1996). doi:10.1016/0550-3213(96)00246-5

[25] R. P. Woodard, *Perturbative quantum gravity comes of age*, Int. J. Mod. Phys. D **23**, 1430020 (2014). doi:10.1142/S0218271814300201

[26] I. Antoniadis, J. Iliopoulos, and T. N. Tomaras, *Quantum instability of de Sitter space*, Phys. Rev. Lett. **56**, 1319 (1986). doi:10.1103/PhysRevLett.56.1319

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

[28] 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