

Bounding the Obstruction to the Modular–Physical Identity

in a Finite Causal Diamond

Quantum-Geometric Correspondence

Marc Sperzel*

Independent Researcher

MSci Physics, King’s College London

July 2, 2026

Abstract

The identity $K = 2\pi H_{\text{phys}}$ — the modular Hamiltonian of an observer’s algebra equals 2π times the physical Hamiltonian — is the Bisognano–Wichmann theorem for the free field, and for gravity it underlies the thermal-time hypothesis, the crossed-product construction of subregion algebras, and a first-order-in- G gravitational decoherence rate $\Gamma_{\text{dec}} = E_G/\hbar$ with $E_G = GM^2/d$. Its proof for gravity faces a specific obstruction: the gravitational dressing is the Newtonian field, supported outside every bounded region, and the flat wedge algebra is Type III₁, admitting no interior factor against which the exterior tail could be localised. We make the obstruction quantitative. For a coherent (Gaussian) dressing the displacement conjugation of the modular Hamiltonian is exact, so the entire deviation $K - 2\pi H_{\text{phys}}$ is carried by a single functional — the boost-weighted energy of the dressing field on the modes outside a finite causal diamond — and we bound it in three steps: off-diagonal (which-path) matrix elements see only the difference of the two branches’ dressings, so the global $1/r$ monopole cancels identically; momentum conservation cancels the dipole, leaving a quadrupole whose exterior weight falls as $(d/h)^4$ with the depth h of the source in the diamond — 10^{-12} for a millimetre superposition one metre inside the laboratory light cone; and the finite diamond, unlike the wedge, satisfies Buchholz–Wichmann energy nuclearity (its observer Hamiltonian is gapped by the $\mathbb{H}^3$ curvature and its level density polynomial, giving the finite index $Z(2\pi) \approx 4.1 \times 10^{-4}$), so the split property supplies the finite localisation constant the wedge lacks. The residual bi-local channel first appears at $\mathcal{O}(G^2)$: its $\mathcal{O}(G)$ part cancels exactly in the Gaussian covariance. The identity is not proven here: the bound is a composition estimate on the named functional in the Gaussian sector, and the two constants we do not compute — the interacting localisation constant and the coefficient of the $\mathcal{O}(G^2)$ residual — are exactly the open content of the conjecture. They are also what a laboratory discrimination between first-order and second-order gravitational decoherence rates, separated by $\sim 3 \times 10^{34}$ at the microgram–millimetre benchmark, would measure.

*ORCID: [0009-0000-6252-3155](https://orcid.org/0009-0000-6252-3155). Email: me@marcsperzel.com

1 Introduction

The identity

$$K = 2\pi H_{\text{phys}} \quad (1)$$

equates the modular Hamiltonian $K = -\ln \Delta$ of an observer’s local algebra with 2π times the physical Hamiltonian the observer’s clock measures. It is the sharp form of the thermal-time hypothesis [16], the operator content behind the crossed-product construction of gravitational subregion algebras [14, 13, 15], and the origin of a first-order-in- G gravitational decoherence rate $\Gamma_{\text{dec}} = E_G/\hbar$ with $E_G = GM^2/d$ [1]. For the free field on a Rindler wedge it is a theorem: Bisognano and Wichmann identified the vacuum modular flow with the boost [4, 5], and the 2π is the same universal factor read off by an accelerated detector [6].

For gravity the identity faces a specific, well-known obstruction. The gravitational dressing of a massive body is its Newtonian field, and the Newtonian field has support outside every bounded region: no strictly local operator creates a mass together with its $1/r$ tail. A dressed state of a finite causal diamond therefore involves modes the diamond does not contain, and the coherent-state argument that identifies the dressed modular Hamiltonian with $2\pi H_{\text{phys}}$ applies only to dressings supported inside the region. Worse, the flat wedge algebra is Type III₁: it admits no interior Type I factor against which the exterior tail could be localised and its effect estimated. The obstruction is usually stated in exactly this qualitative form — the dressing is global, the wedge does not split — and left there, as a reason the identity cannot be proven.

This paper makes the obstruction quantitative. The result is that the entire deviation of the dressed modular Hamiltonian from its vacuum form is carried by a single computable functional — the boost-weighted energy of the dressing field on the modes outside the diamond, the *leakage functional* $W_{\text{ext}}[\alpha]$ of Definition 4.2 — and that in the Gaussian/coherent sector this functional is bounded,

$$W_{\text{ext}}[\alpha_{\text{dressing}}] \leq C \left(\frac{d}{h}\right)^4, \quad C < \infty, \quad (2)$$

for a which-path separation d at depth h inside the diamond. Three mechanisms, each exact or hypothesis-checked in the Gaussian sector, produce the bound. First, the displacement conjugation of the modular Hamiltonian is exact and terminates, so the deviation is linear in the mode displacements and splits cleanly into an interior part — which is $2\pi H_{\text{phys}}$ — and the exterior leakage (Sec. 4). Second, off-diagonal (which-path) matrix elements see only the *difference* of the two branches’ dressings: the global $1/r$ monopole is common-mode and cancels identically, and momentum conservation cancels the resulting dipole, leaving a quadrupole whose exterior weight falls as $(d/h)^4$ — equal to 10^{-12} for a millimetre superposition one metre inside the laboratory light cone (Sec. 4). Third, the finite diamond, unlike the wedge, satisfies the Buchholz–Wichmann energy nuclearity condition — its observer Hamiltonian is bounded below by the curvature gap of $\mathbb{H}^3$ and its level density is polynomial, giving the finite index $Z(2\pi) \approx 4.1 \times 10^{-4}$ — so the split property holds and supplies the finite localisation constant C the wedge lacks (Sec. 5). The residual quadratic (bi-local) channel first appears at $\mathcal{O}(G^2)$: its $\mathcal{O}(G)$ part cancels exactly in the Gaussian covariance (Sec. 6).

What the paper does not do is prove the identity. The bound (2) is a composition estimate on the named functional in the Gaussian sector, not an operator-norm inequality on the interacting algebra, and two numbers are not computed: the interacting value of the localisation constant C , and the interacting coefficient of the $\mathcal{O}(G^2)$ residual. Equation (1) remains a conjecture

(Conjecture 2.3), established at the level of expectation values and in solvable models [1, 12]. The claim of this paper is a reduction: every other candidate source of deviation is fixed — the 2π is universal with no per-mode freedom (Sec. 2), the constraint selects the physical Hamiltonian as the flow’s target (Sec. 3), the exterior leakage is quartically small and the localisation constant finite (Secs. 4–5), and the non-local residual is second order (Sec. 6) — so the conjecture’s fate rests on a sharply isolated, well-posed computation rather than on a diffuse plausibility argument. The same two open constants are what a laboratory discrimination between first-order and second-order gravitational decoherence rates would measure [1, 3].

The paper is organised along the chain of the bound. Section 2 assembles the multimode Bisognano–Wichmann identity $K = 2\pi H_{\text{boost}} + c\mathbb{K}$ and states the gravitational conjecture. Section 3 shows, in a solvable two-sided model, that the Hamiltonian constraint coincides with the modular generator and selects $K_R = f(H_{\text{phys}})$, the KMS state fixing $f = 2\pi(\cdot)$, and that the QED Gauss law does neither. Section 4 derives the exact conjugation, defines the leakage functional, and establishes the monopole cancellation, the multipole ladder, and the $(d/h)^4$ quadrupole scaling. Section 5 verifies energy nuclearity for the diamond and invokes the split property. Section 6 composes the estimate and fixes the $\mathcal{O}(G^2)$ onset of the bi-local residual. Section 7 states what remains open and how it connects to experiment. Appendix A records the diamond kinematics and the observer clock; Appendix B the numerics behind the fitted exponents.

2 The matter backbone: the identity as a multimode operator equation

Before any gravitational dressing enters, the identity $K = 2\pi H_{\text{phys}}$ is an exact theorem of the free field. On the Rindler wedge the Bisognano–Wichmann theorem [4, 5] fixes the modular Hamiltonian of the vacuum to be 2π times the boost generator, and the Unruh construction [6] makes the mechanism explicit: the Minkowski vacuum is a two-mode squeezed state across the horizon, and the squeezing angle carries the 2π . This section assembles that single-frequency fact into a full multimode operator equation, establishes that the coefficient 2π is universal across the spectrum with no per-mode freedom, and states the gravitational target the rest of the paper bounds.

2.1 The multimode Bisognano–Wichmann identity

Theorem 2.1 (Modular backbone of the free field). *Let ϕ be a free scalar field on Minkowski space and $\mathcal{A}(W)$ the algebra of the right Rindler wedge W . Decompose the field into Rindler modes of boost frequency $\omega_i > 0$ with number operators N_i . Then the vacuum modular Hamiltonian $K = -\ln \Delta$ of $\mathcal{A}(W)$ is diagonal in the multimode Fock basis and equals*

$$K = \sum_i \left(2\pi \omega_i N_i - \ln(1 - x_i) \right) = 2\pi H_{\text{boost}} + c\mathbb{K}, \quad x_i = e^{-2\pi\omega_i}, \quad (3)$$

with $H_{\text{boost}} = \sum_i \omega_i N_i$ the boost generator and $c = -\sum_i \ln(1 - x_i)$ a single state-independent c -number. The coefficient 2π multiplies every mode: the per-mode level spacing of K is exactly $2\pi\omega_i$ for all i , with no spectral or per-mode normalisation freedom. The identity holds as an

operator equation on every occupation sector, hence at all moments; in the vacuum,

$$\text{Var}(K) = (2\pi)^2 \text{Var}(H_{\text{boost}}). \quad (4)$$

Equation (3) reads: the generator of modular time on the wedge is the boost energy, calibrated by the universal factor 2π , up to an additive constant that carries only the entanglement entropy. It is the free-field statement of $K = 2\pi H_{\text{phys}}$, exact and mode-by-mode.

Proof. Fix a single Rindler frequency ω . The Minkowski vacuum, restricted to W , is the reduced state of a two-mode squeezed (thermofield-double) state across the horizon with squeezing parameter set by

$$\tanh r = e^{-\pi\omega}, \quad (5)$$

so the wedge sees a thermal state

$$\rho_W = (1-x) \sum_{n \geq 0} x^n |n\rangle\langle n|, \quad x = \tanh^2 r = e^{-2\pi\omega}. \quad (6)$$

The occupation x is fixed by the squeezing (5), not inserted by hand. The single-mode modular Hamiltonian is

$$K_0 = -\ln \rho_W = -\ln(1-x) + (-\ln x) N = -\ln(1-x) + 2\pi\omega N, \quad (7)$$

where the last step uses $-\ln x = 2\pi\omega$ from (6). The level spacing of K_0 is therefore exactly $2\pi\omega$: the modular-to-boost normalisation 2π emerges from the squeezing angle (5), and from nothing else.

The free field is a direct sum over independent Rindler frequencies. The modes do not couple, so the wedge state factorises, $\rho_W = \bigotimes_i \rho_i$, each factor of the form (6) with occupation $x_i = e^{-2\pi\omega_i}$. The logarithm of a tensor product is the sum of logarithms,

$$K = -\ln \rho_W = \sum_i K_i = \sum_i \left(2\pi\omega_i N_i - \ln(1-x_i) \right), \quad (8)$$

which is (3) on collecting the boost generator $H_{\text{boost}} = \sum_i \omega_i N_i$ and the constant $c = -\sum_i \ln(1-x_i)$. Each term contributes the *same* 2π , so $K - 2\pi H_{\text{boost}} = c\mathbb{I}$ is a multiple of the identity: the operator equation holds sector by sector, not merely in expectation. Since the difference is a c -number, the mean of K is shifted by c while the n -th cumulant of K equals $(2\pi)^n$ times the n -th cumulant of H_{boost} for every $n \geq 2$; at $n = 2$ this is the variance identity (4). The universality of 2π across the spectrum is the statement that $-\ln x_i/\omega_i = 2\pi$ independently of ω_i , which is (5) rewritten. $\square$

The content of Theorem 2.1 beyond the textbook single-mode case is that the 2π carries no per-mode knob. A modular-time calibration that differed from mode to mode would smear any coefficient built on it; the free field admits no such freedom, so the one universal $2\pi = \beta_{\text{mod}}$ is the Bisognano–Wichmann value throughout.

2.2 The additive constant and the modular flow

The constant c in (3) is dynamically inert.

Lemma 2.2 (Shift lemma). *For any real c , the operators K and $K + c\mathbb{K}$ generate the same modular flow, and c enters no rate. Only the entropy normalisation depends on it.*

Proof. The modular flow is the one-parameter automorphism group $\sigma_s(A) = \Delta^{is} A \Delta^{-is} = e^{-isK} A e^{isK}$. Replacing $K \rightarrow K + c\mathbb{K}$ multiplies e^{-isK} by the phase e^{-isc} , which commutes with every operator and cancels in the conjugation:

$$e^{-is(K+c)} A e^{is(K+c)} = e^{-isc} e^{-isK} A e^{isK} e^{isc} = e^{-isK} A e^{isK} = \sigma_s(A). \quad (9)$$

The flow, and hence every rate read from it, is unchanged. The constant shifts only $\text{Tr}(\rho K) = S$, the von Neumann entropy, by c . $\square$

Lemma 2.2 isolates the constant c as the entanglement entropy S_0 and removes it from any subsequent bound on a decoherence rate: the gravitational obstruction we bound below is an obstruction to the *operator* content $2\pi H_{\text{boost}}$, never to the additive c .

2.3 Reconstruction of the global generator

The local modular data of Theorem 2.1 carries more than the boost energy: it encodes the full global Hamiltonian. Chen, Lashkari and Leung [12] observe that for the interval $(-R, R)$ in a conformal vacuum the modular Hamiltonian is the local Casini–Huerta–Myers form [11]

$$K(R) = 2\pi \int_{-R}^R \frac{R^2 - x^2}{2R} T_{00}(x) dx, \quad (10)$$

and that the global Hamiltonian $H(R) = \int_{-R}^R T_{00} dx$ is recovered from this purely local datum by a single scale operation,

$$(R \partial_R + 1) K(R) = 2\pi R H(R). \quad (11)$$

The reconstruction is exact for any T_{00} : differentiating (10) in R produces boundary terms carrying the weight $\frac{R^2 - x^2}{2R}|_{x=\pm R} = 0$, which vanish because the modular weight vanishes at the entangling points, so only the weight derivative survives and combines with the $+1$ to give $2\pi R H$. Equation (11) reads: the physical Hamiltonian is not extra information beyond the modular Hamiltonian of a region — it is reconstructible from it, with coefficient 2π . The 2π recovered here is the same 2π of (3) and of the identity this paper studies.

2.4 The gravitational target

Theorem 2.1 and the reconstruction (11) are matter-sector statements: H_{phys} there is the field Hamiltonian on a fixed background. In the gravitational theory H_{phys} is the constraint-reduced Hamiltonian, and the vacuum of the dressed diamond carries a Newtonian field with support outside the diamond. The operator identity is then a conjecture.

Conjecture 2.3 (Modular–physical identity on the dressed diamond). *Let $\mathcal{A}(D)$ be the gravitationally dressed algebra of a finite causal diamond D , K the vacuum modular Hamiltonian of $\mathcal{A}(D)$, and H_{phys} the constraint-reduced Hamiltonian of the linearised gravitational theory.*

Then, as an operator equation on $\mathcal{A}(D)$,

$$K = 2\pi H_{\text{phys}} + \mathcal{O}(G^2). \quad (12)$$

Equation (12) is established at the level of expectation values and in solvable models; as a full interacting operator identity it remains a conjecture [1]. The two correction channels it admits must be kept distinct. The $\mathcal{O}(G^2)$ written in (12) names the *non-local* (bi-local) residual, whose order is fixed in Sec. 6: that channel has no first-order piece. Separately, at finite depth h there is a *local* first-order-in- G term — the exterior leakage of the dressing, which does not vanish identically but is geometrically suppressed as $(d/h)^4$ and vanishes in the deep-source limit $h \rightarrow \infty$; Sec. 4 isolates it and bounds it. The remainder of this paper does not prove (12). It reduces the obstruction to a single computable functional, the boost-weighted energy of the dressing field outside the diamond, and bounds that functional.

3 The constraint selects the modular generator

Theorem 2.1 fixes the free-field modular Hamiltonian, but a generic subregion’s modular Hamiltonian is a bilocal operator bearing no relation to any Hamiltonian. What singles out gravity is that its constraint is not an extra condition layered on the dynamics: the Hamiltonian constraint coincides with the modular generator. In a solvable two-sided model we show that imposing the constraint is imposing modular invariance, that its equilibrium solution forces the one-sided modular Hamiltonian to be 2π times the physical Hamiltonian, and that the spatial Gauss law of electromagnetism — a constraint but not the modular generator — does none of this. This is why gravity, and not electromagnetism, points to the identity of Conjecture 2.3.

3.1 The constraint as the modular generator

Take two sectors, an observer/clock “L” and matter “R”, each with Hamiltonian H and energies E_n . The gravitational Hamiltonian constraint — that the total super-Hamiltonian annihilates physical states [17] — is modelled by

$$\hat{C} |\psi\rangle = 0, \quad \hat{C} = H_L - H_R. \quad (13)$$

The full two-sided modular Hamiltonian of the L|R split at inverse temperature β is $K_{\text{full}} = \beta(H_R - H_L)$ [14]. The two operators therefore coincide up to the temperature factor,

$$\hat{C} = H_L - H_R = -\frac{1}{\beta} K_{\text{full}}. \quad (14)$$

Equation (14) reads: the Hamiltonian constraint and the generator of modular time are the same operator. Imposing $\hat{C} |\psi\rangle = 0$ is identically imposing $K_{\text{full}} |\psi\rangle = 0$, invariance under the full modular flow — the KMS equilibrium condition [16]. The constraint is not imposed on top of the flow; it is the statement that the state is stationary under the flow.

3.2 What the constraint gives and what the temperature gives

The states annihilated by (13) are the energy-correlated thermofield-double class,

$$|\psi\rangle = \sum_n c_n |n\rangle_L |n\rangle_R, \quad (15)$$

for which the reduced matter state $\rho_R = \sum_n |c_n|^2 |n\rangle\langle n|$ is diagonal in the energy basis. Its modular Hamiltonian therefore commutes with the physical Hamiltonian and is a function of it alone:

$$[K_R, H_{\text{phys}}] = 0, \quad K_R = f(H_{\text{phys}}). \quad (16)$$

Equation (16) is the structural content the constraint supplies, and it is the hard part: a generic subregion's modular Hamiltonian is bilocal and unrelated to the dynamics, and the constraint forces it to be a function of the physical Hamiltonian.

What the constraint does not fix is that f is linear with slope 2π . The distinguished modular-invariant (KMS / Hartle–Hawking) state $|c_n|^2 = e^{-\beta E_n} / Z$ gives

$$K_R = -\ln \rho_R = \beta H_{\text{phys}} + \text{const}, \quad \beta = 2\pi \Rightarrow K_R = 2\pi H_{\text{phys}}. \quad (17)$$

The KMS spectrum has uniform gaps equal to 2π . A non-KMS constraint-satisfying state — for instance $|c_n|^2 \propto e^{-\beta E_n - \gamma E_n^2}$ — still satisfies the constraint and still gives $K_R = f(H_{\text{phys}})$, but now with f nonlinear. The linearity and the value 2π are the modular KMS temperature, the Bisognano–Wichmann input of Theorem 2.1, not something the constraint alone delivers.

Proposition 3.1 (Constraint selection in the thermofield model). *In the two-sided thermofield-double model, the Hamiltonian constraint (13) coincides with the modular generator (14), and its imposition forces the one-sided modular Hamiltonian to be a function of the physical Hamiltonian, $K_R = f(H_{\text{phys}})$ with $[K_R, H_{\text{phys}}] = 0$. At the KMS temperature $\beta = 2\pi$ the modular-invariant solution gives $K_R = 2\pi H_{\text{phys}}$; a non-KMS constraint-invariant state gives a nonlinear f .*

Proof. The identity (14) follows from $K_{\text{full}} = \beta(H_R - H_L)$. Any state annihilated by $\hat{C}$ has equal L and R spectral support, hence the Schmidt form (15), so ρ_R is diagonal in the energy basis and $K_R = -\ln \rho_R$ is a function of H_{phys} commuting with it, which is (16). Choosing the modular-invariant weights $|c_n|^2 = e^{-\beta E_n} / Z$ makes $-\ln \rho_R$ affine in E_n with slope β , giving (17); any other constraint-satisfying weight gives a non-affine $-\ln |c_n|^2$, hence nonlinear f . $\square$

Proposition 3.1 is the mechanism: the constraint delivers the structure $K_R = f(H_{\text{phys}})$, and the KMS temperature fixes f to $2\pi H$.

3.3 Why the Gauss law does not select the identity

The contrast with electromagnetism is sharp at the operator level. The QED constraint is the spatial Gauss law $\nabla \cdot E = \rho$, the generator of gauge transformations on a Cauchy slice [18]. It commutes with the boost/modular generator but is not that generator: it is an internal gauge condition, not $H_L - H_R$. Imposing it therefore does not impose modular invariance and does not tie ρ_R to the energy basis. A Gauss-satisfying physical state can entangle the observer with R's gauge index, making ρ_R diagonal in a basis that does not commute with H_R , so that

$$[K_R, H_{\text{phys}}] \neq 0 \quad (\text{Gauss law / QED}). \quad (18)$$

Under the spatial constraint the modular Hamiltonian need not even be a function of the physical Hamiltonian, let alone 2π times it. The distinction is the operator-level content of the statement that gravity has a Hamiltonian constraint, $H_{\text{total}} = 0$, while electromagnetism has only spatial constraints:

	constraint generator	modular Hamiltonian
Gravity	$\hat{C} = H_L - H_R = -K_{\text{full}}/\beta$	$K_R = f(H_{\text{phys}})$; KMS $\Rightarrow 2\pi H_{\text{phys}}$
QED	Gauss law (gauge, $\neq$ modular generator)	$[K_R, H_{\text{phys}}] \neq 0$ in general

The left column is the physical origin of the identity: gravity’s constraint is the modular generator, so equilibrium under it is $K_R \propto H_{\text{phys}}$; electromagnetism’s is not, so no such tie is forced.

Remark 3.1. Two limits of this argument must be stated plainly. First, the 2π is the modular KMS temperature, an input the constraint alone does not deliver: the selection argument forces $K_R \propto H_{\text{phys}}$, and it is the KMS state that fixes the constant to 2π and the function to linear. Second, everything in this section holds in a solvable finite-dimensional thermofield-double model. The field-theoretic realisation on the finite diamond — with the global-support Newtonian dressing whose exterior weight is bounded in Sec. 4 and localised by nuclearity in Sec. 5 — is not carried out here, and $K = 2\pi H_{\text{phys}}$ remains open for the interacting theory. The section exhibits the mechanism by which gravity points to the identity; it does not establish the identity.

4 The error functional and its quadrupole suppression

The identity $K = 2\pi H_{\text{phys}}$ conjectured in Sec. 2 (Conjecture 2.3) equates the modular Hamiltonian of the gravitationally dressed causal diamond with the boost-frame energy of the matter it contains. This section isolates the obstruction to that identity. For a coherent (Gaussian) dressing the deviation $K - 2\pi H_{\text{phys}}$ is computed exactly and shown to reduce to a single scalar functional—the boost-weighted energy of the dressing field on the modes outside the diamond—and that functional is suppressed as $(d/h)^4$ by momentum conservation, with d the which-path separation and h the depth of the source inside the diamond.

4.1 The exact conjugation and the leakage functional

By Bisognano–Wichmann the vacuum modular Hamiltonian of the diamond is the boost generator, diagonal over the wedge modes,

$$K_{\text{vac}} = \sum_k \Omega_k a_k^\dagger a_k, \quad (19)$$

with Ω_k the modular (boost) weight of mode k [4, 5, 7]. A branch carrying its Newtonian field is the field vacuum acted on by a coherent displacement, $|\Phi\rangle = D(\alpha)|0\rangle$ with $D(\alpha) = \exp(\sum_k \alpha_k a_k^\dagger - \alpha_k^* a_k)$, the amplitudes α_k encoding that field. The dressed modular Hamiltonian is $D(\alpha) K_{\text{vac}} D(\alpha)^\dagger$, and this conjugation is exact.

Lemma 4.1 (Exact conjugation of the modular Hamiltonian). *For any coherent displacement $D(\alpha)$,*

$$D(\alpha) K_{\text{vac}} D(\alpha)^\dagger = K_{\text{vac}} - \sum_k \Omega_k (\alpha_k a_k^\dagger + \alpha_k^* a_k) + \sum_k \Omega_k |\alpha_k|^2. \quad (20)$$

The quadratic (modular) part is unchanged; the entire effect of the dressing is the linear term of per-mode magnitude $\Omega_k |\alpha_k|$, the modular weight times the displacement amplitude.

Proof. The displacement acts on the ladder operators by $D(\alpha) a_k D(\alpha)^\dagger = a_k - \alpha_k$ and $D(\alpha) a_k^\dagger D(\alpha)^\dagger = a_k^\dagger - \alpha_k^*$; the Baker–Campbell–Hausdorff series terminates at first order because $[a_k, a_k^\dagger] = 1$ is central. Substituting into each term of Eq. (19),

$$D(\alpha) a_k^\dagger a_k D(\alpha)^\dagger = (a_k^\dagger - \alpha_k^*)(a_k - \alpha_k) = a_k^\dagger a_k - (\alpha_k a_k^\dagger + \alpha_k^* a_k) + |\alpha_k|^2, \quad (21)$$

and summing against Ω_k gives Eq. (20). The number operator $a_k^\dagger a_k$ reappears with its coefficient intact, so the second-moment structure that defines the modular flow is invariant under the displacement; only first moments shift. $\square$

Lemma 4.1 makes the deviation of the dressed modular Hamiltonian from the vacuum one entirely linear in the mode displacements. The coherent-state locality theorem identifies $D(\alpha) K_{\text{vac}} D(\alpha)^\dagger$ with the true modular Hamiltonian of the diamond algebra precisely when $D(\alpha)$ is supported inside the diamond, i.e. when $\alpha_k = 0$ for every mode localized outside; the interior linear term is then $2\pi H_{\text{phys}}$, the boost-frame energy of the dressing field (the same 2π recovered from local modular data in Sec. 2, Theorem 2.1). Splitting the linear term of Eq. (20) by mode support,

$$K - 2\pi H_{\text{phys}} = - \sum_{k \in \text{out}} \Omega_k (\alpha_k a_k^\dagger + \alpha_k^* a_k), \quad (22)$$

identifies the obstruction with the exterior part of the displacement alone: the interior modes reconstruct $2\pi H_{\text{phys}}$ exactly, and every remaining term lives on modes outside the diamond.

Definition 4.2 (Leakage functional). The *leakage functional* of a coherent dressing α is the exterior boost-weighted displacement weight

$$W_{\text{ext}}[\alpha] = \sum_{k \in \text{out}} \Omega_k |\alpha_k| = \int_{\text{outside}} (\text{boost weight}) \times (\text{displacement amplitude}). \quad (23)$$

Equation (22) exhibits $K - 2\pi H_{\text{phys}}$ as a specific operator built from the exterior modes, and $W_{\text{ext}}[\alpha]$ is the scalar that bounds its per-mode content. The bounds of this paper are bounds on $W_{\text{ext}}[\alpha]$: it is a finite functional of the displacement, whereas K and H_{phys} are individually unbounded, so a bare operator norm $\|K - 2\pi H_{\text{phys}}\|$ has no meaning on the wedge algebra without the finite-trace structure supplied in Sec. 5. The functional-level statement—that the deviation is carried entirely by $W_{\text{ext}}[\alpha]$ and vanishes when $W_{\text{ext}}[\alpha] \rightarrow 0$ —is exact. Promoting it to an operator-norm bound on the interacting algebra remains open; that promotion is what Sec. 5 and Sec. 6 address, and it is not established here.

4.2 The relative displacement and the monopole cancellation

The decoherence-relevant matrix elements are off-diagonal in the branch index, $\langle \Phi_L | O | \Phi_R \rangle$ for the modular Hamiltonian and its perturbation. With $|\Phi_A\rangle = D(\alpha_A)|0\rangle$ the branch content enters only through the relative displacement.

Lemma 4.3 (Off-diagonal relative displacement). *The Weyl identity*

$$D(\alpha_L)^\dagger D(\alpha_R) = e^{i \text{Im} \langle \alpha_L, \alpha_R \rangle} D(\alpha_R - \alpha_L) \quad (24)$$

is exact. Consequently every off-diagonal element $\langle \Phi_L | O | \Phi_R \rangle = \langle 0 | D(\alpha_L)^\dagger O D(\alpha_R) | 0 \rangle$ depends on the displacement only through the difference $\alpha_R - \alpha_L$; the common-mode $(\alpha_L + \alpha_R)/2$ cancels.

Proof. For the Heisenberg algebra $[a_k, a_j^\dagger] = \delta_{kj}$ the Baker–Campbell–Hausdorff expansion of $D(\alpha_L)^\dagger D(\alpha_R) = D(-\alpha_L)D(\alpha_R)$ terminates at the commutator, $D(\beta)D(\gamma) = e^{\frac{1}{2}(\beta\gamma^* - \beta^*\gamma)} D(\beta + \gamma)$, giving Eq. (24) with $\beta = -\alpha_L$, $\gamma = \alpha_R$. Writing $\alpha_A = \bar{\alpha} \pm \frac{1}{2}\delta\alpha$ with $\bar{\alpha} = (\alpha_L + \alpha_R)/2$ and $\delta\alpha = \alpha_R - \alpha_L$, the argument of the surviving displacement is $\alpha_R - \alpha_L = \delta\alpha$, independent of $\bar{\alpha}$. $\square$

The field that governs decoherence is therefore the difference $\delta\Phi = \Phi_L - \Phi_R$, not each branch’s monopole. Each branch sources a Newtonian monopole $\Phi_A(x) = -GM/|x - x_A|$ with $|\nabla\Phi_A|^2 \propto r^{-4}$; the difference of two masses separated by d is a dipole,

$$\delta\Phi(x) = -GM\left(\frac{1}{|x - x_L|} - \frac{1}{|x - x_R|}\right) \sim \frac{GMd}{r^2} \quad (r \gg d), \quad (25)$$

with energy density $|\nabla\delta\Phi|^2 \sim d^2/r^6$, two powers more localized than the monopole. The global $1/r$ tail—the piece that lies outside the diamond and would break the conjugation of Lemma 4.1—is common-mode and cancels exactly in the off-diagonal.

The cancellation is quantitative. On a boost-weighted exterior integral over a three-dimensional grid, with the source at depth h inside the diamond, the exterior fraction of the monopole weight is 6.2×10^{-3} at $h = 32$ (in units of the separation), while that of the dipole is 4.3×10^{-6} , smaller by three orders of magnitude at the same depth. Power-law fits give the monopole exterior fraction falling as $h^{-1.6}$ and the dipole as $h^{-3.1}$: the relative-displacement leakage is suppressed by more than one extra power of d/h . The falloff exponents -4 (monopole) and -6 (dipole) are exact and reproduced to machine precision (Appendix B).

4.3 Momentum conservation and the multipole ladder

The dipole difference source of Eq. (25) carries a net dipole moment, which no closed system can support: in any physical preparation the recoil is taken up by the apparatus, and the total mass distribution conserves momentum. Momentum conservation cancels the dipole moment of the difference source, and the leading surviving multipole is the quadrupole,

$$\Phi_{\text{quad}} \sim \frac{GMd^2}{r^3}, \quad |\nabla\Phi_{\text{quad}}|^2 \sim \frac{d^4}{r^8}, \quad (26)$$

two powers more localized than the dipole, four than the monopole. This dipole-to-quadrupole cancellation for a momentum-conserving mass distribution has been derived independently in the influence-functional treatment of gravitational decoherence by Takeda, Tomizuka and Tanaka [20].

The exterior weight follows a clean multipole ladder. With boost weight $|z| \sim r$, volume element $r^2 dr$, and energy density r^{-p} , a multipole of order ℓ has field $\sim r^{-(\ell+1)}$, density exponent $p = 2\ell + 4$, and exterior weight

$$W_{\text{ext}}(\ell, h) = \int_h^R r^{3-p} dr = \int_h^R r^{-(2\ell+1)} dr \propto h^{-2\ell} \quad (\ell \geq 1), \quad (27)$$

so the monopole ($\ell = 0$) gives the marginal logarithm $\ln(R/h)$, while for $\ell \geq 1$ the infrared limit $R \rightarrow \infty$ is finite and the ladder is exact: the dipole ($\ell = 1$) gives h^{-2} and the quadrupole ($\ell = 2$) gives h^{-4} (Sec. 6 takes $R \rightarrow \infty$).

Proposition 4.4 (Quadrupole leakage scaling). *For a momentum-conserving difference source at depth $h \gg a$ inside the diamond, where a is the system–apparatus configuration scale, the leakage functional scales as*

$$W_{\text{ext}}[\alpha] \propto \left(\frac{d}{h}\right)^4. \quad (28)$$

Below the crossover $h \sim a$ the compensating recoil is unresolved and the bare dipole scaling h^{-2} applies.

Proof. Momentum conservation is a single constraint, fixing the $\ell = 1$ moment of the total source but not $\ell = 2$; the generic leading surviving multipole is the quadrupole. The compensating apparatus recoil sits a distance a from the mass, so the difference source is the mass dipole at depth h together with the opposite apparatus dipole at depth $h + a$, of equal and opposite moment but separated by a , leaving a net quadrupole of moment $\sim Mda$. For $h \gg a$ this looks like a point quadrupole and Eq. (27) at $\ell = 2$ gives $W_{\text{ext}} \propto h^{-4}$; with the numerator scale d set by the separation this is $(d/h)^4$. For $h \ll a$ the mass dipole near the boundary dominates and the exterior weight follows the bare dipole h^{-2} ; the steepening onset moves outward with a , so the crossover tracks the configuration scale. $\square$

The exponents are measured directly. On the boost-weighted exterior integral the bare dipole falls as $h^{-2.1}$ and the bare quadrupole as $h^{-4.0}$, and the momentum-conserving composite source falls as $h^{-3.8}$ in the far field ($h \gg a$) and $h^{-1.3}$ in the near field ($h \ll a$), with the crossover tracking a as it is varied (Appendix B). For a laboratory superposition the diamond boundary is far outside any compact apparatus, so the far-field quadrupole scaling holds: a $d = 1$ mm superposition at depth $h = 1$ m gives $(d/h)^4 = 10^{-12}$, and the leakage functional—hence the obstruction to $K = 2\pi H_{\text{phys}}$ on the Gaussian off-diagonal sector—is quartically negligible.

Remark 4.1. The section establishes that for a coherent dressing the deviation $K - 2\pi H_{\text{phys}}$ is carried by the single leakage functional $W_{\text{ext}}[\alpha]$ of Definition 4.2, that W_{ext} sees only the momentum-conserving quadrupole difference field, and that it is suppressed as $(d/h)^4$ for a source deep inside the diamond. It does not establish an operator-norm bound on the interacting algebra: the quadrupole field retains global support, so $W_{\text{ext}}[\alpha]$ is small but nonzero, and the conjugation of Lemma 4.1 is exact only in the Gaussian sector. Promoting the functional bound to a norm bound on the finite-diamond algebra requires the nuclearity estimate of Sec. 5 and the composition of Sec. 6; the identity $K = 2\pi H_{\text{phys}}$ itself remains a conjecture.

5 Nuclearity and the split property of the diamond algebra

The leakage functional of Sec. 4 measures the boost-weighted energy of the dressing on modes outside the diamond, and bounding it requires localising a globally supported dressing operator onto a wedge-internal approximant with a finite localisation constant. The flat Rindler wedge does not supply one: its algebra is Type III₁, which admits no split and hence no intermediate factor onto which the dressing could be projected. The finite diamond does supply one, and this section verifies the phase-space condition that produces it.

5.1 The curvature gap of the diamond observer Hamiltonian

The intermediate factor exists because the observer Hamiltonian of the diamond is bounded below by a strictly positive gap of geometric origin. Under the Casini–Huerta–Myers conformal

map the diamond becomes the ultrastatic cylinder $\mathbb{R} \times \mathbb{H}^3$ [11], on which the modular generator is the cylinder-time Hamiltonian $q = \sum_{\lambda} \omega_{\lambda} a_{\lambda}^{\dagger} a_{\lambda}$ with one-particle energies $\omega_{\lambda} = \sqrt{\lambda}$ set by the eigenvalues λ of the $\mathbb{H}^3$ Laplacian (App. A). The bottom of that spectrum is fixed exactly. In geodesic polar coordinates the s -wave radial Laplace–Beltrami operator on $\mathbb{H}^3$ is $-\Delta f = f'' + 2 \coth(\rho) f'$, and the substitution $f = g / \sinh \rho$ turns it into a constant-potential problem,

$$f'' + 2 \coth(\rho) f' + \lambda f = 0 \iff -g'' + g = \lambda g, \quad (29)$$

an exact identity with no approximation. The reduced operator $-g'' + 1 \cdot g$ on the half-line with $g(0) = 0$ is a free particle in a constant potential $+1$, so its spectrum is the continuum $\lambda = 1 + k^2$, $k \in \mathbb{R}$, and

$$\text{spec}(-\Delta_{\mathbb{H}^3}) = \left[\frac{(n-1)^2}{4}, \infty \right) \Big|_{n=3} = [1, \infty). \quad (30)$$

Equation (30) reads: the diamond observer Hamiltonian is bounded below by the hyperbolic curvature gap $(n-1)^2/4 = 1$, the “ $+1$ ” in Eq. (29) being precisely the curvature term, with no flat-space analogue (where the floor is 0). The one-particle energies therefore satisfy $\omega_{\lambda} = \sqrt{\lambda} \geq 1 > 0$, and the derivation of the bound uses only this gap.

5.2 The nuclearity index

The Buchholz–Wichmann nuclearity condition asks that the map $a \mapsto e^{-\beta q} a |\Omega\rangle$ be nuclear, an index controlled by the partition function $Z(\beta) = \text{Tr } e^{-\beta q}$ [8]. Two features of the diamond make it finite. The spectrum is gapped by Eq. (30), so the integration starts at $\omega = 1$; and the density of states of a three-dimensional Laplacian grows only polynomially — for $\mathbb{H}^3$ the Plancherel density is $\rho(\omega) \sim \omega^2$. A polynomial density damped by the modular weight integrates to a finite value,

$$Z(\beta) = \int_1^{\infty} \rho(\omega) e^{-\beta \omega} d\omega = e^{-\beta} \left(\frac{1}{\beta} + \frac{2}{\beta^2} + \frac{2}{\beta^3} \right) < \infty, \quad \rho(\omega) \sim \omega^2, \quad (31)$$

the closed form holding at gap 1 for every $\beta > 0$, in particular at the modular temperature $\beta_{\text{mod}} = 2\pi$ (App. A), where

$$Z(2\pi) \approx 4.1 \times 10^{-4}. \quad (32)$$

The finiteness is not automatic. A Hagedorn density $\rho \sim e^{\alpha \omega}$ would make the integrand $e^{(\alpha-\beta)\omega}$ and drive $Z(\beta)$ to diverge for $\alpha \geq \beta$; the diamond escapes this exactly because its $\mathbb{H}^3$ /Weyl density is polynomial and its spectrum is gapped. The polynomial degree does not matter — ω^2 and ω^4 both give a finite $Z(2\pi)$ — so the index does not hinge on the precise Plancherel coefficient, only on the sub-exponential growth. Equation (31) reads: the diamond is nuclear because the curvature gap and the polynomial mode count together make the modular weight summable.

5.3 The split property

Nuclearity of the index (31) is a phase-space criterion, and the passage from it to a localising factor is a theorem of Buchholz, D’Antoni and Longo, invoked here rather than reproved.

Theorem 5.1 (Nuclearity implies the split property, after Buchholz–Wichmann and Buchholz–D’Antoni–Longo). *For nested diamonds $D_{\text{in}} \subset D_{\text{out}}$ whose energy nuclearity index $Z(\beta) =$*

$\text{Tr } e^{-\beta q}$ is finite, the inclusion $\mathcal{A}(D_{\text{in}}) \subset \mathcal{A}(D_{\text{out}})$ splits: there exists a Type I factor $\mathcal{N}$ with

$$\mathcal{A}(D_{\text{in}}) \subset \mathcal{N} \subset \mathcal{A}(D_{\text{out}}), \quad (33)$$

and the split distance, the deviation of the split from an exact tensor factorisation, is controlled by the nuclearity index [8, 9, 10].

Theorem 5.1 is not our result; the implication nuclearity $\Rightarrow$ split is standard operator algebra. Our contribution is the verification of its hypothesis for the gravitationally relevant object: the diamond observer Hamiltonian is gapped by Eq. (30) and its density is polynomial, so the index (31) is finite and the theorem applies. The matter-sector instance of the same argument — the L^2 -nuclearity of the ball inclusion $B(e^{-2\pi s}R) \subset B(R)$ — has been established directly by Chen, Lashkari and Leung [12]. The Type I factor $\mathcal{N}$ of Eq. (33) is the wedge-internal approximant: the globally supported dressing has a representative inside $\mathcal{N}$ with an error equal to the split distance, and that error is finite because the index is. This is the localisation constant the leakage bound of Sec. 6 requires, absent on the Type III₁ wedge and present on the diamond.

5.4 The interacting extension

The construction so far is free-field: the density $\rho \sim \omega^2$ is the Gaussian mode count. The localisation constant remains finite once interactions are turned on, provided the level density stays sub-exponential.

Proposition 5.2 (Finiteness of the localisation constant under interactions). *Assume that, below the Hagedorn/Planck threshold, the interacting graviton self-coupling corrects the diamond level density $\rho(\omega) \sim \omega^2$ only by polynomial terms, $\rho(\omega) = \omega^2 + c_3\omega^3 + c_4\omega^4 + \dots$ with finitely many terms. Then the nuclearity index remains finite,*

$$Z(\beta) = \int_1^\infty (\omega^2 + c_3\omega^3 + c_4\omega^4 + \dots) e^{-\beta\omega} d\omega < \infty, \quad (34)$$

the split property persists, and the localisation constant stays finite, rescaled by an $\mathcal{O}(1)$ factor relative to its free value.

Proof. Each monomial ω^p damped by $e^{-\beta\omega}$ integrates to $Z_p(\beta) = e^{-\beta} P_p(1/\beta)$ with P_p a polynomial of degree $p+1$, finite for every finite p and every $\beta > 0$. A finite sum of finite terms is finite, so Eq. (34) converges at $\beta_{\text{mod}} = 2\pi$ regardless of the degree or the coefficients. Nuclearity therefore holds, the split of Theorem 5.1 persists, and the ratio $Z_{\text{inter}}/Z_{\text{free}}$ is bounded — an $\mathcal{O}(1)$ rescaling of the constant. The only way to break this is a genuine Hagedorn density $\rho \sim e^{\alpha\omega}$ with $\alpha \geq \beta$, for which Z diverges; that is a Planck-scale string effect outside the effective-field-theory regime in which the identity is posed. $\square$

The premise of Proposition 5.2 — polynomial corrections to the level density below the Hagedorn/Planck threshold — is an assumption, stated as such, not a computed fact. What the proposition establishes is the conditional: given sub-Hagedorn growth, the localisation constant is finite throughout the effective-field-theory regime, which is the regime of the laboratory prediction. The explicit interacting *value* of the constant requires the full interacting modular data and is not computed here; it is one of the two open numbers isolated in Sec. 6.

6 The composition estimate and the order of the residual

The three preceding sections each fix one input to the obstruction. The constraint selection of Sec. 3 fixes the target of the identity to be $2\pi H_{\text{phys}}$. The leakage functional of Sec. 4 isolates the only deviation as the exterior boost-weighted energy of the dressing and shows it is suppressed as $(d/h)^4$ for a momentum-conserving source. The nuclearity of Sec. 5 supplies a finite localisation constant. This section composes them into a single bound on the leakage functional, and separately establishes the order at which the genuinely non-local part of the obstruction first appears.

6.1 The Gaussian-sector composition estimate

Proposition 6.1 (Gaussian-sector composition estimate). *In the Gaussian/coherent sector, for a momentum-conserving (quadrupole) source at depth $h \gg a$ inside a diamond of separation scale d , the leakage functional $W_{\text{ext}}[\alpha_{\text{dressing}}]$ of Sec. 4 (Def. 4.2) is bounded by*

$$W_{\text{ext}}[\alpha_{\text{dressing}}] \leq C (d/h)^4, \quad C < \infty, \quad (35)$$

with the constant C finite by the nuclearity of the diamond (Theorem 5.1, Proposition 5.2), and the non-local (quadratic) sector contributing to the obstruction only at $\mathcal{O}(G^2)$.

Proof. The constraint selection (Prop. 3.1) fixes the target $2\pi H_{\text{phys}}$ and identifies the sole first-order deviation as the leakage functional $W_{\text{ext}}[\alpha_{\text{dressing}}]$, the boost-weighted displacement weight on modes outside the diamond. For a source obeying momentum conservation the leading monopole and dipole weights cancel, the first surviving multipole is the quadrupole, and the exterior radial integral of its boost-weighted profile is the ladder of Prop. 4.4,

$$W_{\text{ext}}(h) = \int_h^\infty r^{3-p} dr = \frac{h^{4-p}}{p-4} \Big|_{p=8} = \frac{1}{4} h^{-4}, \quad (36)$$

the quadrupole weight $p = 8$ giving the h^{-4} falloff. Restoring the separation scale, $W_{\text{ext}} \leq C (d/h)^4$. The constant C is the split distance of the wedge-internal approximant: the Type I factor of Theorem 5.1 localises the globally supported dressing with an error equal to that distance, which is finite because the nuclearity index (31) is finite, and which stays finite under interactions by Proposition 5.2. The quadratic sector contributes through the inter-region covariance, which is $\mathcal{O}(G^2)$ by Proposition 6.2 below and vanishes for coherent states, so it enters the obstruction only at second order. $\square$

Equation (35) bounds the named leakage functional $W_{\text{ext}}[\alpha_{\text{dressing}}]$ of Def. 4.2, not a bare operator norm $\|K - 2\pi H_{\text{phys}}\|$. The distinction is not cosmetic: K and H_{phys} are unbounded, so a norm on their difference is not defined on the interacting algebra, and the leakage functional is the well-posed object the constraint analysis actually produces. Promoting the functional estimate (35) to an operator-norm inequality on the interacting Type II algebra — the passage from a bound on the exterior weight to a bound on the operator itself — is not derived here; it is one of the two open constants named at the end of this section. What Eq. (35) establishes is that the exterior leakage, the entire first-order content of the obstruction, vanishes as $(d/h)^4$ with a finite coefficient.

The deep-source limit reads the identity off the bound. As $h \rightarrow \infty$ at fixed d the right-hand side of Eq. (35) vanishes, so the identity $K = 2\pi H_{\text{phys}}$ holds exactly in the deep-source limit of the Gaussian/coherent sector — the qualifier being essential: it is the Gaussian sector in which every step, from the Bisognano–Wichmann backbone through the coherent dressing to the free-field split, is under control. At laboratory depth $h/d \sim 10^3$ the geometric channel (36) is smaller than 10^{-10} , negligible, which leaves the $\mathcal{O}(G^2)$ non-Gaussian channel as the sole surviving residual.

6.2 The order of the non-local residual

The obstruction has a second channel beyond the exterior leakage: a bi-local term in K coupling two points inside the diamond, which would spoil the first-order rate were it $\mathcal{O}(G)$. It is not.

Proposition 6.2 (The bi-local residual is $\mathcal{O}(G^2)$). *Write $K = 2\pi H_{\text{phys}} + \delta K_{\text{nonlocal}}$ with $\delta K_{\text{nonlocal}}$ the bi-local part. In the Gaussian covariance model, the $\mathcal{O}(G)$ dressing is a first-moment (coherent displacement) shift that leaves the covariance — and hence the quadratic part of K — untouched, so $\delta K_{\text{nonlocal}}$ carries no $\mathcal{O}(G)$ term. The bi-local seed is the inter-region covariance*

$$\sigma_{12} = \frac{1}{2}(\sigma_{ss} - \sigma_{aa}) \sim g^2, \quad (37)$$

even in the coupling g , so that

$$\delta K_{\text{nonlocal}} = \mathcal{O}(\sigma_{12}) = \mathcal{O}(G^2), \quad \text{no } \mathcal{O}(G) \text{ piece.} \quad (38)$$

Proof. For a Gaussian state the modular Hamiltonian is fixed entirely by the covariance; first moments do not enter. The Newtonian dressing of a branch is a coherent displacement D : it shifts $\langle x \rangle$ at $\mathcal{O}(G)$, the field being $\Phi \propto GM/r$, but leaves the covariance unchanged. So K , and in particular its bi-local part, is untouched at first order, $\delta K_{\text{nonlocal}} = 0$ at $\mathcal{O}(G)$. A bi-local term requires a non-zero inter-region covariance σ_{12} : for a Gaussian state K is block-diagonal, hence purely local, if and only if the covariance is, and the bi-local coefficient is analytic in σ_{12} . That covariance is a second-moment effect. In the solvable model of two region modes sharing one bath, each coupling with strength g — the analogue of two mass elements sourcing one gravitational field — the symmetric/antisymmetric split gives

$$\sigma_{ss} = \frac{1}{4}[(w^2 + \sqrt{2}g)^{-1/2} + (w^2 - \sqrt{2}g)^{-1/2}], \quad \sigma_{aa} = \frac{1}{2}(w^2)^{-1/2}, \quad (39)$$

so that $\sigma_{12} = \frac{1}{2}(\sigma_{ss} - \sigma_{aa})$ vanishes at $g = 0$, is even in g — the odd, first-order contribution cancelling exactly by the $1 \leftrightarrow 2$ symmetry — and grows as g^2 , doubling g quadrupling σ_{12} . Hence $\sigma_{12} \sim g^2$, Eq. (37), and $\delta K_{\text{nonlocal}} = \mathcal{O}(\sigma_{12}) = \mathcal{O}(G^2)$ with no $\mathcal{O}(G)$ piece, Eq. (38). $\square$

The graviton reading of this counting is that a one-graviton dressing is a first moment and produces no bi-local term, while a two-graviton correlation is a covariance and seeds one at $\mathcal{O}(G^2)$: one power of G per moment, two moments for a bi-local term. This is a counting argument at the level of the Gaussian covariance model, not a field-theory computation; the interacting value of $\delta K_{\text{nonlocal}}$ for genuine interacting gravitons is the second of the two open numbers, and Eq. (38) fixes only its order, not its coefficient.

Remark 6.1 (The residual content of the conjecture). Everything in the chain that fixes the *structure* of the obstruction is in place: the constraint fixes the target $2\pi H_{\text{phys}}$; the leakage functional carries the entire first-order deviation and vanishes as $(d/h)^4$; the nuclearity of the diamond makes the localisation constant finite; and the non-local residual first appears at $\mathcal{O}(G^2)$. What is not computed is two numbers — the interacting value of the localisation constant C in Eq. (35), and the interacting coefficient of the $\mathcal{O}(G^2)$ residual in Eq. (38). These two constants are exactly the open content of the conjecture (12): the identity is reduced to them, and to nothing else.

7 Discussion

7.1 What is fixed and what is open

The chain of Secs. 2–6 reduces the conjecture $K = 2\pi H_{\text{phys}} + \mathcal{O}(G^2)$ to two numbers. Fixed: the 2π is the universal Bisognano–Wichmann calibration with no per-mode freedom (Theorem 2.1); the additive constant is dynamically inert (Lemma 2.2); the constraint selects H_{phys} as the target of the flow (Prop. 3.1); the entire first-order deviation is the leakage functional, suppressed as $(d/h)^4$ by the monopole cancellation and momentum conservation (Prop. 4.4); the localisation constant is finite by nuclearity, and stays finite under sub-Hagedorn interactions (Theorem 5.1, Prop. 5.2); the bi-local residual first appears at $\mathcal{O}(G^2)$ (Prop. 6.2). Open: the interacting value of the localisation constant C in Eq. (35), and the interacting coefficient of the $\mathcal{O}(G^2)$ residual — equivalently, the promotion of the functional estimate to an operator-norm statement on the interacting Type II algebra. These are interacting quantum-gravity quantities; no step in this paper computes them, and the identity remains a conjecture exactly as stated in Conjecture 2.3.

Other candidate sources of coefficient freedom in the decoherence rate built on the identity are closed elsewhere and enter no part of the present bound: the entropy-constant residues of the diamond construction — the trace-anomaly normalisation and the entangling-surface edge modes — are additive c-numbers decoupled from the flow by Lemma 2.2 and are treated in the companion graviton paper [2], and the residual window in the rate coefficient, $C_{\text{rate}} \in [2/\pi, 1]$, is an observable convention pinned in the canonical-core paper [1]. With those channels closed, the coefficient freedom of the rate and the operator content of the identity coincide: both live in the two constants above.

7.2 The experiment that measures the residual

The two open constants are not of purely formal interest: they are what the laboratory would measure. The identity at first order gives the decoherence rate $\Gamma_{\text{dec}} = E_G/\hbar$, linear in G ; perturbative graviton exchange with an unconstrained environment gives a rate at $\mathcal{O}(G^2)$ [1, 20]. The expansion parameter separating them at the canonical benchmark — a microgram superposition at millimetre separation — is $\epsilon_G = GM/(c^2 d) = 7.4 \times 10^{-34}$, and the ratio of the first-order to the second-order rate is $(M_P/M)^2 (d/\ell_P) \approx 3 \times 10^{34}$: thirty-four orders of magnitude between the conjectured first-order rate ($\tau_{\text{dec}} = \hbar d/(GM^2) = 1.6$ ns at the benchmark) and the second-order rate. A which-path experiment at picogram masses that resolves the gravitational decoherence floor reads off the order in G directly, and with it the fate of the operator equation [1, 3]. The spectral shape of the same decoherence channel — the

modular-thermal knee at $f = E_G/(2\pi\hbar)$ — supplies the finer discriminant between the modular mechanism and Markovian collapse models of the Diósi–Penrose class [21, 22, 3].

7.3 Relation to nearby results

Three strands of recent work touch the same objects. The crossed-product constructions [14, 13, 15] build Type II subregion algebras by adjoining an observer with a bounded-below Hamiltonian; the diamond observer of Appendix A, gapped by the $\mathbb{H}^3$ curvature (Sec. 5.1), is precisely such an observer, and the finite nuclearity index of Sec. 5 is the phase-space input those constructions presuppose. Chen, Lashkari and Leung [12] established the matter-sector instance of the programme rigorously — the global Hamiltonian reconstructed from local modular data with the same 2π , and L^2 -nuclearity for the conformal ball inclusion — which anchors the free-field end of the present chain; the gravitational, dressed case is what the bound of Sec. 6 addresses. Danielson, Satishchandran and Wald [19] showed that a black-hole or cosmological horizon decoheres superpositions through the soft radiation the two branches emit through the horizon, the number of entangling emissions growing linearly in the time the superposition is held open; the mechanism here is the intrinsic analogue — the observer’s own diamond horizon at the modular temperature $\beta_{\text{mod}} = 2\pi$, no external horizon required — and the momentum-conservation quadrupole structure of Sec. 4.3 matches the multipole counting of their emission argument, derived independently in the influence-functional setting by Takeda, Tomizuka and Tanaka [20].

7.4 Outlook

The obstruction to $K = 2\pi H_{\text{phys}}$ was a qualitative no-go: the dressing is global, the wedge does not split. It is now an estimate with a stated remainder — the leakage is quartically small in d/h , the localisation constant is finite, the non-local channel is second order — and a two-constant residual that is well-posed, isolated, and shared with the observable the experiment targets. Computing either constant in interacting quantum gravity, or measuring the order in G of gravitational decoherence in the laboratory, decides the conjecture. Nothing else does.

A Diamond kinematics and the observer clock

The modular temperature $\beta_{\text{mod}} = 2\pi$ and the ultrastatic $\mathbb{R} \times \mathbb{H}^3$ frame used in Secs. 5 and 6 are fixed by the diamond’s own geometry, not imported by analogy. This appendix records the kinematic facts that fix them. The modular flow of a finite causal diamond $D = \{|t| + r < R\}$ is generated by the conformal Killing vector [7]

$$\xi = \frac{\pi}{R} \left[\frac{R^2 - t^2 - r^2}{2} \partial_t - t r \partial_r \right], \quad (40)$$

which reads: the diamond’s vacuum flows under a boost-like conformal symmetry that fixes the diamond and its light cone.

The observer worldline. The diamond has isometry group $\text{SO}(3)$ of rotations about its centre. The unique timelike worldline left pointwise fixed by that group is the central geodesic $r = 0$, $|t| < R$ — a straight inertial line in Minkowski space, the diamond’s intrinsic analogue of

the de Sitter static-patch worldline, selected by symmetry alone. It is an integral curve of the modular generator (40): at $r = 0$,

$$\xi^r|_{r=0} = -\frac{\pi}{R} t r|_{r=0} = 0, \quad \xi^t|_{r=0} = \frac{\pi}{2R}(R^2 - t^2), \quad (41)$$

so ξ is purely timelike on the centre and the flow carries the observer along their own worldline.

Modular time and the horizon. Integrating the flow $dt/d\sigma = \xi^t|_{r=0} = (\pi/2R)(R^2 - t^2)$ with $t(0) = 0$ gives

$$t(\sigma) = R \tanh\left(\frac{\pi\sigma}{2}\right), \quad \sigma \in (-\infty, \infty) \Rightarrow t \in (-R, R), \quad (42)$$

so the observer reaches the future and past tips $t = \pm R$ only as $\sigma \rightarrow \pm\infty$: the diamond tips are a modular horizon at infinite modular time, the diamond crossed in finite proper time $\int d\tau = 2R$ but in unbounded modular parameter — the hallmark of a boost-like modular flow on a finite region.

The uniform cylinder clock. Under the Casini–Huerta–Myers Weyl rescaling $\hat{g} = \Omega^{-2}\eta$ the diamond maps to the ultrastatic cylinder $\mathbb{R} \times \mathbb{H}^3$ [11], on which ξ becomes a constant-norm Killing field, $|\xi|_{\hat{g}}^2 = -\pi^2 R^2/4$ everywhere. The $\hat{g}$ -proper time along the central curve is therefore linear in the modular parameter,

$$\hat{\tau}(\sigma) = \frac{\pi R}{2} \sigma, \quad \frac{d\hat{\tau}}{d\sigma} = \sqrt{-|\xi|_{\hat{g}}^2} = \frac{\pi R}{2} = \text{const}, \quad (43)$$

a clock ticking uniformly in modular time, exactly the uniform clock the crossed product needs; its conjugate energy is the observer Hamiltonian q whose spectrum is bounded below by the $\mathbb{H}^3$ curvature gap of Sec. 5.1.

The modular temperature. The modular flow is KMS at the normalisation-free inverse temperature

$$\beta_{\text{mod}} = 2\pi, \quad (44)$$

with position-dependent Tolman temperature $T_{\text{loc}}(r) = 1/(2\pi f(r))$, $f(r) = (R^2 - r^2)/(2R)$; at the centre $T_{\text{loc}}(0) = 1/(\pi R)$, redshifting to the uniform cylinder value $T_{\text{cyl}} = 1/(\pi R)$. Equation (44) is the value used throughout the paper: the observer carries a thermal clock at the universal modular temperature 2π . The associated crossed-product refinements — the trace-anomaly normalisation of the entropy constant and the entangling-surface edge modes — are not part of the present bound; they are treated elsewhere [2].

B Exterior-weight numerics

This appendix records the grid construction and the fitted power laws behind the leakage suppression of Sec. 4, specifically the monopole–dipole cancellation of Sec. 4.2 (Lemma 4.3) and the multipole ladder of Sec. 4.3 (Proposition 4.4).

B.1 The boost-weighted exterior integral

The diamond is modelled by the Rindler wedge $z > 0$, its complement the left wedge $z < 0$, with the entangling plane at $z = 0$. A source of separation d is placed at depth h inside the wedge (distance h from the entangling plane). The obstruction of Definition 4.2 is the boost-weighted weight of the displacement field outside the wedge, so the discretized quantity is

$$W = \int |z| |\nabla \Phi|^2 d^3x, \quad \frac{W_{\text{ext}}}{W_{\text{int}}} = \frac{W(z < 0)}{W(z > 0)}, \quad (45)$$

with the boost weight $|z|$ standing for the modular weight Ω_k of Eq. (19). The field Φ is the monopole $-GM/r$, the dipole difference of Eq. (25), or the momentum-conserving composite of Eq. (26); the integral of Eq. (45) is evaluated on a three-dimensional grid and the exterior fraction $W_{\text{ext}}/W_{\text{int}}$ is read off. The falloff exponents are checked separately on the radial proxy $W_{\text{ext}}(p, h) = \int_h^R r^{3-p} dr$ of Eq. (27), integrated by Simpson's rule to $R = 10^4$.

B.2 The depth sweep and the monopole–dipole cancellation

Table 1 gives the exterior fraction of Eq. (45) as the source depth h is swept, for the monopole (each branch's Newtonian field) and the dipole (the relative displacement of Lemma 4.3), in units where the separation is unity.

depth h	monopole $W_{\text{ext}}/W_{\text{int}}$	dipole $W_{\text{ext}}/W_{\text{int}}$	ratio (dipole/monopole)
4	1.7×10^{-1}	2.6×10^{-3}	1.5×10^{-2}
8	6.8×10^{-2}	3.4×10^{-4}	5.0×10^{-3}
16	2.3×10^{-2}	4.1×10^{-5}	1.8×10^{-3}
32	6.2×10^{-3}	4.3×10^{-6}	7.0×10^{-4}

Table 1. *Boost-weighted exterior fraction versus source depth. The dipole (relative displacement) exterior weight is smaller than the monopole at every depth and falls faster; the ratio decreases monotonically toward zero, so the leakage of the decoherence-relevant field is suppressed with depth. Power-law fits: monopole $\sim h^{-1.6}$, dipole $\sim h^{-3.1}$.*

The dipole exterior fraction is smaller than the monopole at every depth and falls by more than one extra power of h ; the common-mode $1/r$ monopole tail, which carries the global support that would break the conjugation of Lemma 4.1, has cancelled in the difference.

B.3 The multipole ladder and the crossover

Table 2 gives the fitted exterior-weight exponents for the pure multipoles and for the momentum-conserving composite of Proposition 4.4. The pure-multipole values confirm the ladder $W_{\text{ext}}(\ell, h) \propto h^{-2\ell}$ of Eq. (27): the exterior exponent equals -2ℓ for $\ell = 1, 2, 3$ on the radial proxy, and the boost-weighted grid integral reproduces the dipole -2 and quadrupole -4 to within fit error.

The composite source, built as the mass dipole at depth h and the opposite apparatus dipole at depth $h + a$, reproduces the two regimes of Proposition 4.4: the far-field exponent -3.8 approaches the pure quadrupole -4.0 for $h \gg a$, the near-field exponent -1.3 tracks the bare dipole for $h \ll a$, and increasing a pushes the crossover outward. For a laboratory geometry $h \gg a$, so the quadrupole scaling of Eq. (28) governs the leakage.

source	fitted exterior-weight exponent
bare dipole ($\ell = 1$)	-2.1
bare quadrupole ($\ell = 2$)	-4.0
momentum-conserving composite, far field ($h \gg a$)	-3.8
momentum-conserving composite, near field ($h \ll a$)	-1.3

Table 2. Fitted power-law exponents of the boost-weighted exterior weight. The bare dipole and quadrupole match the ladder exponents -2ℓ . The momentum-conserving composite (mass dipole plus compensating apparatus dipole separated by the configuration scale a) crosses over from the near-field dipole h^{-2} to the far-field quadrupole h^{-4} ; the steepening onset moves outward as a is increased, confirming that the crossover tracks a .

References

- [1] M. Sperzel, *The Quantum-Geometric Correspondence: Three Axioms and Gravitational Decoherence at Order G* , (2026), Quantum-Geometric Correspondence, canonical core paper.
- [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*, (2026), Quantum-Geometric Correspondence series.
- [3] M. Sperzel, *Modular Spectroscopy: A Spectral Fingerprint of Gravitationally-Induced Decoherence*, (2026), Quantum-Geometric Correspondence series.
- [4] J. J. Bisognano and E. H. Wichmann, *On the duality condition for a Hermitian scalar field*, J. Math. Phys. **16**, 985 (1975). [doi:10.1063/1.522605](https://doi.org/10.1063/1.522605)
- [5] 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](https://doi.org/10.1063/1.522898)
- [6] W. G. Unruh, *Notes on black-hole evaporation*, Phys. Rev. D **14**, 870 (1976). [doi:10.1103/PhysRevD.14.870](https://doi.org/10.1103/PhysRevD.14.870)
- [7] P. D. Hislop and R. Longo, *Modular structure of the local algebras associated with the free massless scalar field theory*, Commun. Math. Phys. **84**, 71 (1982). [doi:10.1007/BF01208372](https://doi.org/10.1007/BF01208372)
- [8] D. Buchholz and E. H. Wichmann, *Causal independence and the energy-level density of states in local quantum field theory*, Commun. Math. Phys. **106**, 321 (1986). [doi:10.1007/BF01454978](https://doi.org/10.1007/BF01454978)
- [9] D. Buchholz, C. D’Antoni, and R. Longo, *Nuclear maps and modular structures II: Applications to quantum field theory*, Commun. Math. Phys. **129**, 115 (1990). [doi:10.1007/BF02096782](https://doi.org/10.1007/BF02096782)
- [10] S. Doplicher and R. Longo, *Standard and split inclusions of von Neumann algebras*, Invent. Math. **75**, 493 (1984). [doi:10.1007/BF01388641](https://doi.org/10.1007/BF01388641)
- [11] 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](https://doi.org/10.1007/JHEP05(2011)036)
- [12] Y. Chen, N. Lashkari, and K. L. Leung, *Local approximations of global Hamiltonian from inclusion of algebras*, arXiv:2512.25062 (2026). [arXiv:2512.25062](https://arxiv.org/abs/2512.25062)

- [13] V. Chandrasekaran, R. Longo, G. Penington, and E. Witten, *An algebra of observables for de Sitter space*, J. High Energy Phys. **02** (2023) 082, arXiv:2206.10780. [doi:10.1007/JHEP02\(2023\)082](https://doi.org/10.1007/JHEP02(2023)082)
- [14] E. Witten, *Gravity and the crossed product*, J. High Energy Phys. **10** (2022) 008, arXiv:2112.12828. [doi:10.1007/JHEP10\(2022\)008](https://doi.org/10.1007/JHEP10(2022)008)
- [15] K. Jensen, J. Sorce, and A. J. Speranza, *Generalized entropy for general subregions in quantum gravity*, J. High Energy Phys. **12** (2023) 020, arXiv:2306.01837. [doi:10.1007/JHEP12\(2023\)020](https://doi.org/10.1007/JHEP12(2023)020)
- [16] A. Connes and C. Rovelli, *Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories*, Class. Quantum Grav. **11**, 2899 (1994), arXiv:gr-qc/9406019. [doi:10.1088/0264-9381/11/12/007](https://doi.org/10.1088/0264-9381/11/12/007)
- [17] B. S. DeWitt, *Quantum theory of gravity. I. The canonical theory*, Phys. Rev. **160**, 1113 (1967). [doi:10.1103/PhysRev.160.1113](https://doi.org/10.1103/PhysRev.160.1113)
- [18] W. Donnelly, *Decomposition of entanglement entropy in lattice gauge theory*, Phys. Rev. D **85**, 085004 (2012), arXiv:1109.0036. [doi:10.1103/PhysRevD.85.085004](https://doi.org/10.1103/PhysRevD.85.085004)
- [19] D. L. Danielson, G. Satishchandran, and R. M. Wald, *Black holes decohere quantum superpositions*, Int. J. Mod. Phys. D **31**, 2241003 (2022), arXiv:2205.06279. [doi:10.1142/S0218271822410036](https://doi.org/10.1142/S0218271822410036)
- [20] H. Takeda, S. Tomizuka, and T. Tanaka, *Gravitationally Induced Quantum Decoherence of Macroscopic Objects*, arXiv:2606.04099 (2026). [arXiv:2606.04099](https://arxiv.org/abs/2606.04099)
- [21] 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](https://doi.org/10.1016/0375-9601(87)90681-5)
- [22] R. Penrose, *On gravity's role in quantum state reduction*, Gen. Relat. Gravit. **28**, 581 (1996). [doi:10.1007/BF02105068](https://doi.org/10.1007/BF02105068)