The price of a classical force channel, measured end-to-end in a lattice field theory:

no force from decoherence alone, saturation of the Kafri–Taylor–Milburn bound, and apparatus-independence of the heating floor

The Emergent Gravity Project*

July 2026 — draft for comment, not yet peer reviewed

Abstract

We study, in an exactly solvable Gaussian testbed — a one-dimensional lattice scalar field with steady states obtained from Lyapunov equations and dynamics from exact diagonalization of the drift, no Monte Carlo anywhere — whether and at what cost a decohering medium can *produce* or *carry* a force between two objects. Three results. (i) **Decoherence alone produces no attraction.** Record-creating elastic scatterers ("quantum Le Sage grains") repel weakly, the repulsion is proportional to the medium's bulk absorption and extrapolates to zero in the transparent limit, and a real thermal bath only makes it more repulsive; the *same* scattering operator, made static (coherent), attracts with a Casimir-type force that our stress observable reproduces against the exact ground-state energy to 0.5–1%. Maxwell's nineteenth-century cancellation argument against elastic shadow gravity survives quantization and decoherence. (ii) **A genuine coupling can be carried classically, at a measured price.** Two slow (below-band) masses, delocalized as in the Bose–Marletto–Vedral protocol and coupled linearly to the field, entangle through the vacuum ($E_N \sim 1.2$, long range); local monitoring of the mediating field kills the entanglement at a finite threshold while the transmitted force is *exactly* unchanged (first moments decouple from position dephasing; verified to five digits). At the entanglement-death threshold, the rate at which the environment acquires records about the masses' positions saturates the Kafri–Taylor–Milburn lower bound $\Lambda = K/2\hbar$ to within 2% across all separations tested. (iii) **The heating floor is apparatus-independent.** Replacing pointwise monitors by smeared probes of width R changes the required monitoring rate sixfold but leaves the sources' heating at threshold flat to $\sim 3\%$ — an information–disturbance equality. Consequently the spatial resolution R_0 of Diósi–Penrose-type models cannot be fixed by any minimal-heating principle: it is structurally free, and only experiment (underground spontaneous-radiation bounds from below, short-range gravity tests from above) constrains it.

1 Motivation

If gravity is mediated by a quantum field, two masses interacting only gravitationally can become entangled; proposed Bose–Marletto–Vedral (BMV) experiments aim to detect exactly this [1, 2]. The alternative is that gravity is a *classical channel*: the environment continuously acquires information about where energy is located, and the force is carried by that classical record, as in

*Contact via emergent-gravity.com. Numerical work performed in human–AI collaboration (Claude, Anthropic); all results are reproducible from the public scripts cited in Sec. 7.

the measurement-and-feedback models of Kafri, Taylor and Milburn (KTM) [3] and their field-theoretic completion by Tilloy and Diósi [4]. Classical channels cannot create entanglement, so they predict a null BMV result — but they cannot be free: KTM showed that reproducing a force constant K through a classical channel requires decohering the sources at a rate at least $\Lambda \geq K/2\hbar$, which for Newtonian gravity reproduces the (regularized) Diósi–Penrose (DP) decoherence rate [5, 6]. The parameter-free DP model is already excluded by the Gran Sasso underground search for the spontaneous radiation that such decoherence implies [7], leaving a finite corridor for the model's resolution parameter R_0, bounded below by radiation searches and above by short-range tests of the inverse-square law [8].

This note does three things in a single exactly solvable arena. First, it closes a loophole older than the quantum discussion: could the force itself *come from* the decoherence — a quantum revival of Le Sage's 1748 shadow gravity [9], with vacuum modes in place of corpuscles and which-path records in place of absorption? The classical theory dies by Poincaré's heating estimate and by Maxwell's cancellation theorem for elastic scattering [10]; the only conceivable quantum repair is that what attenuates in the shadow is *coherence*, which is not conserved and hence need not be re-emitted. This is a sharp, computable claim, and we computed it (Sec. 3): it is false. Second, having established that monitoring cannot create a force, we measure what it *can* do: render an existing coupling classical, and we exhibit the KTM bound as an equality at the entanglement-death threshold, end-to-end in an extended field with real propagation rather than in a two-oscillator abstraction (Sec. 4). Third, we ask whether any economy principle selects the monitoring resolution, and find a clean negative with a simple mechanism (Sec. 5). Section 6 states limitations and what we believe is genuinely new.

2 The testbed

The medium is a harmonic chain (lattice scalar field) of N sites, $H_f = \frac{1}{2}\sum_i p_i^2 + \frac{1}{2}\sum_i (x_{i+1}-x_i)^2 + \frac{\mu^2}{2}\sum_i x_i^2$, with a small mass μ regulating the zero mode ($\hbar = c = 1$; $N = 256$, $\mu = 0.02$ in Sec. 3; $N = 192$ field sites, $\mu = 0.05$ afterwards). All states are Gaussian and all dynamics are generated by drift matrices A and diffusion matrices D acting on the $2N \times 2N$ covariance σ: steady states solve the Lyapunov equation $A\sigma + \sigma A^\top + D = 0$; time evolution is $\sigma(t) = e^{At}(\sigma_0 - \sigma_\infty)e^{A^\top t} + \sigma_\infty$ with e^{At} from exact diagonalization. Nothing is sampled.

Two technical devices matter. *(a) Vacuum-continuation baths.* With site-dependent damping η_i (weak in the bulk, ramped at the edges) and the diffusion chosen as $D = \frac{1}{2}(E\sigma_{\rm gs} + \sigma_{\rm gs}E)$, where $E = \mathrm{diag}(\eta, \eta)$ and $\sigma_{\rm gs}$ is the exact ground-state covariance, the undriven steady state *is* the exact ground state (we verify $\max|\sigma_{\rm NESS} - \sigma_{\rm gs}| = 7 \times 10^{-13}$): the finite chain behaves as a window on an infinite vacuum, and outgoing radiation is absorbed without preparing artifacts. *(b) A validated force observable.* The force on an object is read from the discontinuity of the momentum-flux T^{xx} across it, with a four-configuration (both/only-A/only-B/neither) subtraction that cancels single-object and boundary systematics. Calibration: for two *static* quadratic defects the observable must reproduce the exact Casimir force $-dE_0/dd$ computed from ground-state energies. It does, at ratios 0.995, 0.997, 1.002, 1.010 for separations $d = 12, 16, 24, 32$; a single noisy grain shows a parasitic force below one part in 10^7 of its radiated stress.

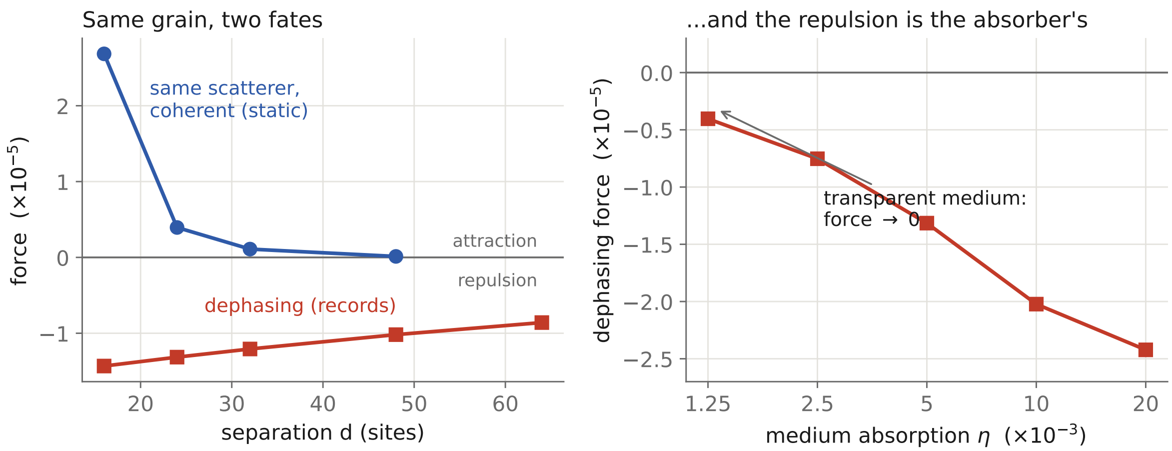

Figure 1: **The shadow kill-gate.** Left: the same scattering operator attracts when static (coherent; a Casimir force, fast-decaying) and produces a weak *repulsion* when noisy (record-creating), at all separations. Right: the dephasing repulsion at $d = 24$ is proportional to the medium's bulk absorption η and extrapolates to approximately zero in the transparent limit ($\lesssim 5 \times 10^{-7}$, within systematics; stable against $N = 256 \to 384$ at 0.3%).

3 Result 1: decoherence does not attract

A "grain" is a fluctuating elastic scatterer: coupling $\xi(t)\,(c_a^\top x)^2/2$, where c_a is the discrete gradient of a narrow Gaussian centred at site a (an index fluctuation: infrared-safe, band-edge-safe) and ξ is white noise of strength γ_a. Averaged over noise, each grain dephases the modes crossing it — it creates which-path records — while absorbing nothing on average. This is the strongest Gaussian implementation of the "coherence shadow" repair of Le Sage gravity: if tagged modes pushed less, two grains would attract with the characteristic $\gamma_A\gamma_B$ two-record scaling.

The verdict (Fig. 1): the interaction is *repulsive* at every separation ($F = -1.4 \times 10^{-5}$ to -0.9×10^{-5} for $d = 16$–64 at $\gamma = 10^{-2}$), scales as $\gamma_A\gamma_B$ in the perturbative regime, is proportional to the bulk absorption η and vanishes with it, and becomes *more* repulsive in a real thermal bath ($F = -2.8\times10^{-6}$ at $T = 0$ to -3.8×10^{-5} at $T = 2$, $\gamma = 5\times10^{-3}$): there is no shadow to cast even when the bath is real. The control experiment is the sharpest statement: the identical coupling operator, made static, attracts ($+2.7 \times 10^{-5}$ at $d = 16$, decaying steeply). **Decoherence does not create the attraction; it destroys it.** Physically, Maxwell's cancellation survives: the wave a grain rediffuses with scrambled phase carries exactly the momentum it removed from the coherent beam, so a coherence deficit exerts no pressure; the only real effect of dephasing a $T = 0$ field is its unavoidable parametric backaction, whose radiation, absorbed by the medium, supplies the small repulsion. We registered the criteria (sign, scaling, d- and η-dependence, thermal behaviour) before the runs; the repair — and with it every Le Sage mechanism, classical or quantum — fails its own gate. Caveats: one dimension, Gaussian dynamics, white noise, two coupling choices tested (site and strain; both agree); Maxwell's argument is dimension-blind.

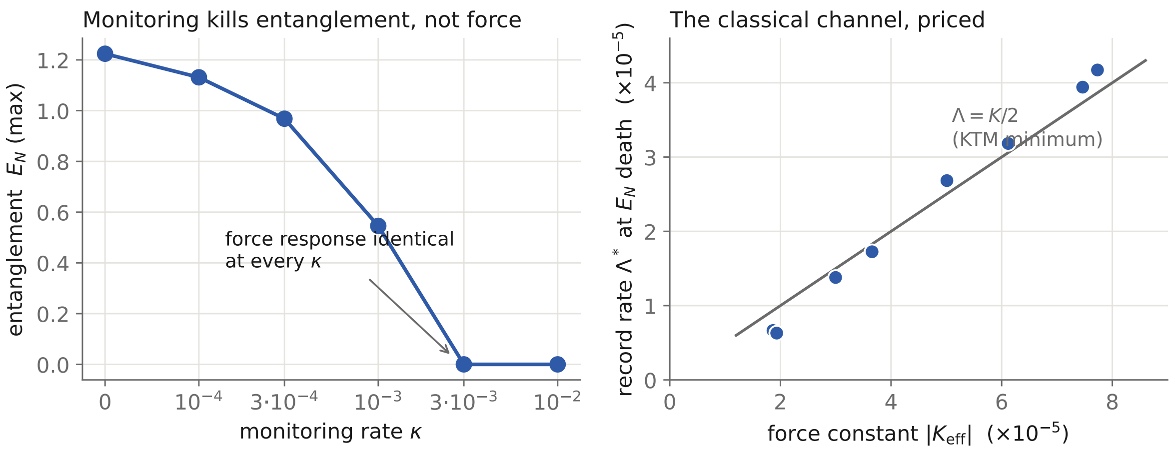

Figure 2: **The monitored channel.** Left: mediated entanglement between the delocalized masses ($d = 12$) dies at finite κ; the transmitted force is identical at every κ. Right: across eight configurations (four separations $\times$ three couplings), the record rate about the masses at the entanglement-death threshold lies on the KTM line $\Lambda = K_{\rm eff}/2$.

4 Result 2: the classical channel, priced

We now give the channel a genuine coupling to carry. Two "masses" (oscillators of frequency $\Omega_m = 0.04$, far *below* the field band $[\mu, 2]$: no radiative leakage, and the field follows adiabatically — the Newtonian regime of slow sources and a fast medium) couple linearly to the field, $g\, x_A\, (c_a^\top x)$, with g calibrated for a 40% frequency softening. The static channel is long-ranged: the effective force constant $K_{\rm eff} = g^2\, c_a^\top K_f^{-1} c_b$ decays on the field's correlation length. The masses are delocalized as in BMV: initial x-squeezing $r = 2$ (position spread $\times e^{2r} \approx 55$), initial state a product. The "vacuum that keeps records" is a position-dephasing Lindblad $\sqrt{\kappa}\, x_i$ on every field site in a window covering the channel.

Proposition 1 (force blindness)**.** *For jump operators linear in the position quadratures, the first-moment equations of motion contain no κ: any force read from mean response is exactly independent of the monitoring. (Immediate from $[x_i, [x_i, x_j]] = [x_i, [x_i, p_j]] = 0$.) Verified numerically: the response of mass A to a displaced mass B is identical to all printed digits ($\langle x_A \rangle = -0.018104$) from $\kappa = 0$ to $\kappa = 0.3$.*

Unmonitored, the channel entangles: $E_N^{\max} = 1.24, 1.22, 1.15, 1.07$ at $d = 8, 12, 16, 20$ — long range, tracking $K_{\rm eff}$. Monitored, the entanglement dies at a finite threshold κ^* (bisected to $\max_t E_N < 10^{-3}$ over $t \le 600$) at the cost of a few percent extra heating. The price is the point: the rate at which the environment acquires information about a mass's position is $\Lambda_A = \kappa \sum_{i\in\rm win} \big(g\, [K_f^{-1} c_a]_i\big)^2$ (the static imprint of x_A in the field, read at rate κ; the adiabatic regime makes this exact). At threshold, for the fiducial series $d = 8, 12, 16, 20$:

$$\Lambda^*/|K_{\rm eff}| \;=\; 0.539,\ 0.528,\ 0.520,\ 0.536,$$

i.e. the KTM bound $\Lambda \ge K/2$ is *saturated* within 2% (Fig. 2, right). Reduced couplings drift slightly below (0.33–0.47), as expected: our death criterion (one initial state, finite horizon)

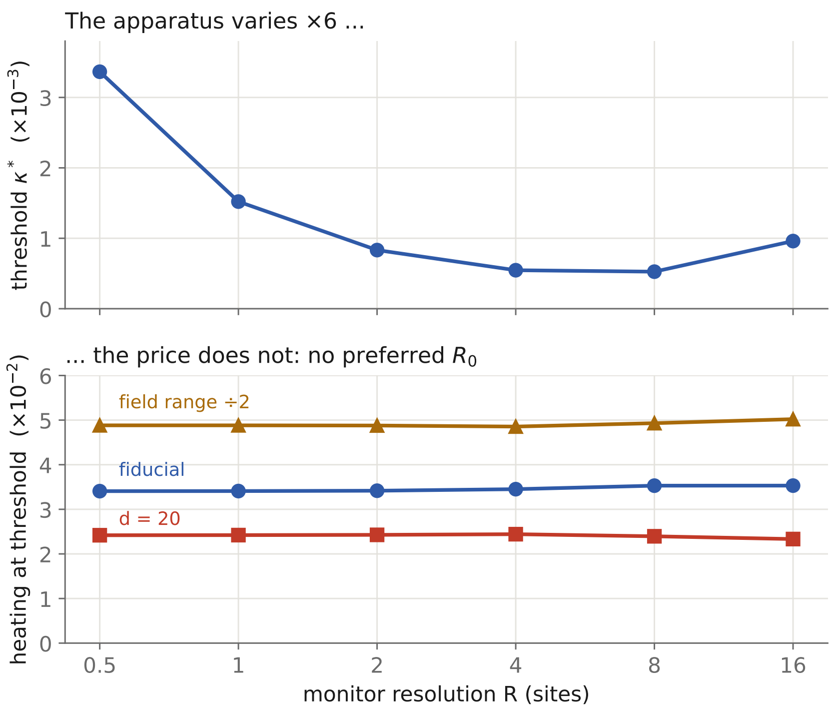

Figure 3: **No preferred resolution.** Top: the threshold monitoring rate $\kappa^*(R)$ varies sixfold with the probe width. Bottom: the heating of the sources at threshold is flat to $\sim 3\%$ in all three configurations (and the residual slope changes sign between them).

is weaker than classicality for all states, so the strict bound need not bind. In real units, $K = 2Gm^2/d^3$ turns this floor into the regularized DP rate: *intact force, null BMV, DP-floor decoherence* is one channel theorem, not three separate predictions, and the toy exhibits all three limbs in a single run.

5 Result 3: the resolution is structurally free

DP-type models carry a resolution parameter R_0 — the scale at which the environment resolves positions — and one might hope to *derive* it: a "lazy vacuum" would monitor at whatever resolution minimizes the heating it must inflict at fixed classicality. We replaced the pointwise monitors by Gaussian probes of width R (grid spacing 2 sites; the per-probe normalization drops out of all threshold quantities), bisected $\kappa^*(R)$ for $R = 0.5$–16, and measured the masses' heating there, in three configurations (fiducial; field correlation length halved; $d = 20$).

The result (Fig. 3): while κ^* varies by $\times 6$, the heating at threshold is flat to $\sim 3\%$, with a residual slope that changes sign between configurations. There is no optimum.

Observation 1 (information–disturbance equality). *For linear monitoring, adiabatic elimination of the field leaves on each mass a single induced channel with jump operator $\propto x_A$: one coefficient Λ controls both the record rate (position decoherence) and the momentum diffusion $D_{pp} = \hbar^2\Lambda$ (heating). Since the classicality threshold pins $\Lambda^* = K/2\hbar$ (Sec. 4), the source heating at threshold is fixed by K alone, independent of the apparatus. The measured Λ^*/K stays* *near $\frac{1}{2}$ across the whole R scan (*0.534*–*0.449*), confirming the mechanism.*

Two consequences. R_0 cannot be derived from economy: within this class it is a genuinely free parameter, and only the experimental corridor constrains it — from below, $R_0 \geq 0.54\,\text{Å}$ (Gran Sasso spontaneous-radiation bound [7]); from above, $R_0 \lesssim 40\,\mu\text{m}$ (Eöt-Wash inverse-square tests [8]); in between, levitated-superposition experiments probe the plateau law $\Gamma \simeq Gm^2/\hbar R_0$ for $\Delta x > R_0$. Conversely, the *radiation floor on the sources is apparatus-independent*: underground searches test the bound coefficient itself, not a modelling choice.

6 Discussion

What we believe is new here: (i) a computational refutation of the decoherence-shadow (quantum Le Sage) mechanism, with pre-registered criteria and a coherent-versus-decohering control on the identical operator — Maxwell's theorem extended, empirically, to quantum records; (ii) an end-to-end demonstration, in an extended field with real propagation and a BMV-style protocol, that the KTM decoherence bound is *saturated* at the entanglement-death threshold, tying together force-blindness, BMV-nullity and the DP floor as limbs of one theorem; (iii) the apparatus-independence of the threshold heating and its corollary, that R_0 is structurally underivable within the class. We are not aware of prior numerical work exhibiting (ii) or (iii) in a field-theoretic (rather than two-oscillator) setting; we welcome pointers to prior art and will credit them in place, as this project has done before.

Limitations, stated plainly: one spatial dimension; Gaussian states and dynamics throughout; position-basis monitoring; white noise; the entanglement-death criterion uses one (squeezed) initial state and a finite horizon; the KTM coefficient is measured at $\pm 2\%$ on the fiducial series and $\pm 35\%$ across all configurations. None of these touches the negative results' logic — Maxwell's cancellation and the first-moment decoupling are structural — but the quantitative saturation deserves an analytic proof in this setting, which we do not have.

7 Reproducibility

Every number and figure comes from three public scripts (plain `numpy`/`scipy`; exact linear algebra, no sampling): `n13_dephasing_shadow.py` (∼1 min), `n14_classical_channel.py` (∼2 min), `n15_record_resolution.py` (∼10 min). The full chronicle of the exercise, including the errors kept on the record, is at emergent-gravity.com/lesage.

References

[1] S. Bose *et al.*, *Spin entanglement witness for quantum gravity*, Phys. Rev. Lett. **119**, 240401 (2017).

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

[3] D. Kafri, J. M. Taylor and G. J. Milburn, *A classical channel model for gravitational decoherence*, New J. Phys. **16**, 065020 (2014).

[4] A. Tilloy and L. Diósi, *Sourcing semiclassical gravity from spontaneously localized quantum matter*, Phys. Rev. D **93**, 024026 (2016).

[5] L. Diósi, *A universal master equation for the gravitational violation of quantum mechanics*, Phys. Lett. A **120**, 377 (1987).

[6] R. Penrose, *On gravity's role in quantum state reduction*, Gen. Relativ. Gravit. **28**, 581 (1996).

[7] S. Donadi *et al.*, *Underground test of gravity-related wave function collapse*, Nat. Phys. **17**, 74 (2021).

[8] J. G. Lee *et al.*, *New test of the gravitational* $1/r^2$ *law at separations down to 52* μm, Phys. Rev. Lett. **124**, 101101 (2020).

[9] G.-L. Le Sage, *Lucrèce Newtonien*, Mémoires de l'Académie Royale des Sciences et Belles-Lettres de Berlin (1784).

[10] J. C. Maxwell, *Atom*, Encyclopædia Britannica, 9th ed. (1875) — the elastic-collision cancellation argument against Le Sage's mechanism.

[11] P. W. Milonni, R. J. Cook and M. E. Goggin, *Radiation pressure from the vacuum: physical interpretation of the Casimir force*, Phys. Rev. A **38**, 1671 (1988).