DATED HISTORY

The Mercury Rung: A Substrate Test for Spin-2 Gravity in QATNU

May design for a substrate test. The completed finite-window search and source-response checks are documented in the consolidated null result.

Manuscript · p. 1, research program

Spin-2 search · null result

Evidence · Negative QATNU spin-2 search

The Mercury Rung

A Substrate Test for Spin-2 Gravity in QATNU

Joshua Farrow
Independent Research

May 2026

Manuscript status. This is a research-program paper. It defines a narrow target for an emergent-gravity model and gives mathematical and numerical criteria by which the target can succeed or fail. It does not claim a completed theory of quantum gravity, a detected new force, or a replacement for general relativity.

Abstract

A large class of emergent-gravity proposals begins with a substrate more primitive than smooth spacetime and tries to recover the familiar gravitational field as a long-wavelength phenomenon. The central difficulty is not producing a suggestive analogy with geometry. The central difficulty is deriving the dynamical content of gravity: a massless helicity-2 field, governed at weak field by the Pauli-Fierz structure, coupled universally to conserved sources, and yielding quantitative outputs not inserted by hand.

This paper formulates a test for one candidate substrate, called the Quantum Automaton Tensor-Network Universe (QATNU), built from local quantum degrees of freedom on vertices and dynamical bond registers on edges. The model naturally supports a scalar timing sector: an internal scalar load Λ_i can renormalize local clock rates through $\omega_{\text{eff}} = \omega/(1 + \alpha\Lambda_i)$, and a weighted graph Laplacian can organize scalar response, diffusion, and effective dimension. That scalar sector is useful, but it is not yet gravity. Scalar timing is holonomy-free, does not carry transverse-traceless tensor modes, and cannot by itself reproduce the full weak-field content of general relativity.

The paper therefore draws a sharp boundary. A scalar timing sector may be an observable projection of the substrate, while the real discovery test is whether QATNU can autonomously generate a spin-2 sector without projecting it in. The proposed target is concrete: two gapless transverse-traceless helicity modes; a low-momentum dispersion $\Omega^2 = c_g^2 k^2 + O(k^4)$; a $1/k^2$ long-distance propagator; Pauli-Fierz kinetic ratios; Ward-identity-like decoupling of nonphysical polarizations; and a universal coupling to a conserved effective source. The model must also confront the Weinberg-Witten obstruction: any successful route must explain why the hypotheses of the theorem do not apply, rather than silently assuming a composite graviton can appear in an ordinary Lorentz-covariant quantum field theory.

This version adds three necessary hardening layers. First, the tensor question is framed as a criticality question: tensor anisotropy must be subcritical, critical, or runaway under coarse-graining, and only the critical case can support a massless long-wavelength sector. Second, the numerical protocol is pre-registered: the Hamiltonian, observables, graph family, coarse-graining map, fitting windows, and thresholds must be fixed before spectra are inspected. Third, the local-chart problem is made explicit: edge directions used to define tensor observables must not

secretly insert the geometry the calculation claims to derive. The primary output target remains the dimensionless ratio

$$R_g = \frac{c_g}{c_{\text{clock}}},$$

with secondary outputs such as stiffness ratios and critical exponents treated as diagnostics rather than substitutes for the main test.

The title refers to the methodological role of Mercury's anomalous perihelion precession in general relativity: not because the historical situations are equivalent, but because a serious theory must eventually produce a number that can refuse it. If the QATNU substrate yields the spin-2 sector under these criteria, the proposal becomes a genuine candidate for emergent gravitational physics. If it does not, then the scalar timing framework remains a controlled internal-medium or projection theory, but not a derivation of gravity.

Contents

1 The problem **4**

1.1 Why a scalar clock sector is not enough 4

2 The Mercury criterion **5**

3 The candidate substrate: QATNU **5**

3.1 What lives where 7

4 The scalar timing sector **7**

4.1 Weighted graph closure 8

4.2 Dimension from the same operator 8

4.3 A severe scalar benchmark 9

5 Why scalar timing is holonomy-free **9**

6 The textbook ladder and the projection trap **10**

7 The Pauli-Fierz target **10**

8 No-go constraints and the required escape hatch **11**

9 Where a tensor could come from in QATNU **12**

9.1 The local chart problem 12

9.2 Hydrodynamic logic 13

10 Criticality rather than plausibility **13**

11 The numerical spin-2 protocol **14**
11.1 Pre-registration before simulation 14
11.2 Step 1: choose graph families 15
11.3 Step 2: define substrate observables before projection 15
11.4 Step 3: measure spectra 15
11.5 Step 4: fit the infrared propagator 16
11.6 Step 5: test source coupling 16
11.7 Step 6: extract the primary Mercury number 16
11.8 Finite-size and infrared diagnostics 17

12 Phase-0 null controls and minimal failed substrates **17**

13 Analytical route: fast-edge elimination **19**

14 Internal fields versus external signals **19**

15 What the framework solves **20**
15.1 It separates projection from derivation 20
15.2 It gives the model a hard target 20
15.3 It makes failure informative 20
15.4 It identifies the first possible discovery point 20

16 What would count as success **21**

17 What would count as failure **21**

18 Discussion **22**

19 Conclusion **22**

A Notation registry **23**

B Minimal implementation sketch **24**

C Scalar closure benchmark appendix **24**

D References **25**

1 The problem

General relativity describes gravity not as a scalar force but as the dynamics of spacetime geometry. In weak fields, this geometry can be written as a perturbation of flat spacetime,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \qquad |h_{\mu\nu}| \ll 1. \tag{1}$$

The propagating gravitational degrees of freedom are not arbitrary components of $h_{\mu\nu}$. After constraints and gauge redundancy are accounted for, the radiative sector contains two transverse-traceless helicity-2 modes. Any substrate theory that hopes to recover gravity must explain why these are the surviving long-wavelength modes.

This is the core difficulty for emergent gravity. It is comparatively easy to build a model that contains clocks, local phases, diffusion, graph distances, entanglement proxies, or redshift-like scalar responses. It is much harder to show that the same model produces the tensor dynamics of gravity. A scalar redshift law can resemble one face of weak-field gravity, but gravity is not exhausted by redshift. Gravity also contains gravitomagnetic structure, gravitational radiation, constraint equations, gauge redundancy, and universal coupling to stress-energy.

The purpose of this paper is to separate these layers cleanly. It introduces a concrete substrate model, identifies the scalar timing sector it naturally supports, states why that sector is insufficient, and defines the exact spin-2 test that would convert the model from an elegant projection language into a candidate theory of emergent gravity.

1.1 Why a scalar clock sector is not enough

A scalar timing field has the form

$$\frac{\mathrm{d}\tau_i}{\mathrm{d}t} = \frac{\omega_{\mathrm{eff},i}}{\omega} = \frac{1}{1+\alpha\Lambda_i} \simeq 1 - \alpha\Lambda_i, \qquad |\alpha\Lambda_i| \ll 1, \tag{2}$$

where Λ_i is an internal scalar load and α is a clock susceptibility. This kind of law is physically meaningful. It says that a local clock can be slowed by a local internal load. In a weak static gravitational field, a clock also slows according to

$$\frac{\mathrm{d}\tau}{\mathrm{d}t} \simeq 1 + \frac{\Phi}{c^2}, \tag{3}$$

where Φ is the Newtonian potential. This analogy makes scalar timing worth studying.

But the analogy has a limit. A scalar potential has no curl, no transverse-traceless radiation, and no spin-2 polarization structure. If one writes a scalar as an exact form on a graph or manifold, loop sums telescope. The scalar field can describe local clock lag, but it cannot by itself describe the tensor sector of gravity.

The central claim of this paper is therefore deliberately narrow:

> The scalar timing sector may be the first clock-readable projection of a structured substrate, but the decisive gravitational test is autonomous spin-2 emergence.

2 The Mercury criterion

A theory becomes more than a language when it produces an output that was not already put into it. The historical benchmark is not that general relativity was beautiful; it is that the field equations produced Mercury's anomalous perihelion precession as a consequence of the theory's structure. The number was not merely fitted by adding a correction term. It emerged from equations that were also constrained by the equivalence principle, special relativity, and the Newtonian limit.

This paper uses that history methodologically. The goal is not to compare personalities or eras. The goal is to identify the kind of target that prevents a research program from endlessly reframing itself. For a substrate theory of gravity, the corresponding target is:

> Derive the weak-field spin-2 gravitational sector from the substrate, including a nontrivial numerical output, without inserting the spin-2 sector into the observable map by hand.

This is the Mercury criterion. It has three parts.

1. **The target must be known enough to test.** The weak-field limit of gravity has a clear benchmark: a massless helicity-2 field with Pauli-Fierz structure and universal source coupling.
2. **The target must not be inserted.** A transverse-traceless projection may be used as a diagnostic after the fact, but the construction of the candidate field cannot simply define h^{TT} and then rediscover it.
3. **The target must output a number.** A qualitative resemblance is not enough. The substrate must predict a dimensionless quantity or scaling law that was not fixed in the assumptions.

The number should be dimensionless whenever possible. Dimensionful scales can be moved by units, matching conventions, or phenomenological anchoring. Dimensionless ratios are harder to hide. This is why constants such as the electromagnetic fine-structure constant play such an important methodological role in physics. QATNU does not yet derive charge, photons, or the electron, so the fine-structure constant is not its first target. The analogous first target inside the present model is instead

$$R_g = \frac{c_g}{c_{\mathrm{clock}}}, \tag{4}$$

where c_g is extracted from the tensor sector and c_{clock} from the scalar/causal-clock sector in the same microscopic units. If R_g is chosen by convention, it is not an output. If it is independently extracted and stabilizes under scaling, it is the first real number the substrate has produced.

The rest of the paper formalizes this criterion for QATNU.

3 The candidate substrate: QATNU

QATNU stands for Quantum Automaton Tensor-Network Universe. The name denotes a family of local quantum automaton models in which matter-like degrees of freedom live on graph vertices and entanglement-carrying bond registers live on graph edges. The specific dynamical rule used here is called Self-Referential Quantum Information Dynamics (SRQID): local matter configurations promote or demote nearby bond registers, and those bond registers feed back into effective clock behavior.

The model is intentionally finite and discrete. Let $G = (V, E)$ be a finite graph. Each vertex $i \in V$ carries a two-level matter system with Hilbert space $\mathcal{H}_i^{\mathrm{mat}} \cong \mathbb{C}^2$. Each edge $e = \langle i, j \rangle \in E$ carries a finite bond register with Hilbert space $\mathcal{H}_{ij}^{\mathrm{bond}} \cong \mathbb{C}^{\chi_{\max}}$. The total Hilbert space is

$$\mathcal{H} = \bigotimes_{i \in V} \mathcal{H}_i^{\mathrm{mat}} \otimes \bigotimes_{\langle i,j \rangle \in E} \mathcal{H}_{ij}^{\mathrm{bond}}. \tag{5}$$

A minimal Hamiltonian is

$$H = H_{\mathrm{mat}} + H_{\mathrm{bond}} + H_{\mathrm{int}}, \tag{6}$$

with

$$H_{\mathrm{mat}} = \frac{\omega}{2} \sum_{i \in V} X_i, \tag{7}$$

$$H_{\mathrm{bond}} = \sum_{\langle i,j \rangle} \sum_{m=0}^{\chi_{\max}-1} \epsilon_m \Pi_{ij}^{(m)}, \tag{8}$$

$$H_{\mathrm{int}} = \lambda \sum_{\langle i,j \rangle} F_{ij} \otimes (R_{ij} + R_{ij}^{\dagger}). \tag{9}$$

Here X_i and Z_i are Pauli operators on vertex i, $\Pi_{ij}^{(m)} = |m\rangle\langle m|_{ij}$ projects the bond register onto rung m, and

$$F_{ij} = \frac{1}{2}(I - Z_i Z_j) \tag{10}$$

is a local frustration projector. The ladder operator

$$R_{ij} = \sum_{m=0}^{\chi_{\max}-2} \sqrt{m+1}\, |m+1\rangle\langle m|_{ij} \tag{11}$$

promotes the bond register. The parameter λ controls how strongly matter frustration promotes or demotes bond structure.

Because H is a finite sum of local bounded terms, the dynamics

$$U(t) = e^{-iHt} \tag{12}$$

is unitary and has a finite Lieb-Robinson velocity on finite-degree graphs. The model is therefore not a collapse model, not an explicitly nonlocal signalling model, and not a continuum field theory assumed at the outset.

For a graph with $N = |V|$ vertices and $|E|$ edges, the Hilbert-space dimension is

$$\dim \mathcal{H} = 2^N \chi_{\max}^{|E|}. \tag{13}$$

This is not a cosmetic detail. It sets the computational ceiling. Exact diagonalization is useful for validation, null controls, and very small lattices; larger scaling tests require Krylov time evolution, tensor-network methods, symmetry reduction, or Monte Carlo sampling of correlators.

3.1 What lives where

The model has three levels that must not be confused.

1. **Microscopic level.** The microscopic state is a vector or density matrix on $\mathcal{H}$. This level contains matter qubits, bond-register amplitudes, phases, correlations, and all microscopic entanglement information.
2. **Internal effective level.** Coarse observables such as bond occupation, scalar load, weighted graph structure, and local clock rate live on the internal graph. These are not yet external spacetime fields.
3. **External observational level.** An external sensor at position x can only see the internal graph through a coupling kernel. For a scalar observable, the most general linear readout is

$$\Lambda_{\text{ext}}(x) = \sum_i C(x,i)\Lambda_i. \tag{14}$$

   If $C(x,i)$ is negligible, the scalar load remains an internal medium variable. If $C(x,i)$ is short-ranged, the observable is near-field. If $C(x,i)$ has a power-law tail, the scalar could look field-like outside the medium. These are empirical distinctions, not assumptions.

The same caution applies to tensor emergence. A tensor-like internal observable is not automatically a gravitational field in external spacetime. It must propagate, couple, and scale correctly.

4 The scalar timing sector

The most immediate coarse observable in QATNU is bond occupation. Let

$$n_{ij} = \sum_{m=0}^{\chi_{\max}-1} m\Pi_{ij}^{(m)} \tag{15}$$

be the rung-number operator on edge $\langle i,j\rangle$. A local scalar load at vertex i can be defined by

$$\Lambda_i = \sum_{j\in\partial i} \log(1+\langle n_{ij}\rangle), \tag{16}$$

or, in the dilute regime, $\Lambda_i \simeq \sum_{j\in\partial i}\langle n_{ij}\rangle$.

The clock law is then

$$\omega_{\text{eff},i} = \frac{\omega}{1+\alpha\Lambda_i}, \tag{17}$$

where ω is the bare matter-clock frequency and α is an effective susceptibility. In weak load,

$$\frac{\delta\omega_i}{\omega} \equiv \frac{\omega_{\text{eff},i}-\omega}{\omega} \simeq -\alpha\Lambda_i. \tag{18}$$

This is a legitimate internal timing law. It can be probed in finite simulations by preparing states, evolving under H, and extracting oscillation frequencies from local observables such as $\langle X_i(t)\rangle$ or $\langle Y_i(t)\rangle$.

4.1 Weighted graph closure

A promoted graph can be represented by a weighted adjacency matrix $W = (w_{ij})$ with degree matrix

$$D_{ii} = \sum_j w_{ij} \tag{19}$$

and weighted Laplacian

$$L_w = D - W. \tag{20}$$

For a function f on vertices,

$$(L_w f)_i = \sum_j w_{ij}(f_i - f_j). \tag{21}$$

A leading scalar closure has the screened graph-Poisson form

$$(\kappa L_w + \mu^2 I)\Lambda = A\rho, \tag{22}$$

where ρ is an effective source, κ is graph stiffness, μ is a screening scale, and A is a source amplitude. A minimal dynamical completion is

$$\partial_t \Lambda + \Gamma(\kappa L_w + \mu^2 I)\Lambda = \Gamma A \rho(t). \tag{23}$$

If $L_w \phi_n = \lambda_n \phi_n$, then the relaxation rates obey

$$\tau_n^{-1} = \Gamma(\kappa \lambda_n + \mu^2). \tag{24}$$

This is experimentally and numerically useful: transients should track graph-Laplacian eigenvalues if the scalar closure is real.

4.2 Dimension from the same operator

The same Laplacian also defines a diffusion semigroup

$$K(t) = e^{-tL_w} \tag{25}$$

and a mean return probability

$$P_{\text{return}}(t) = \frac{1}{|V|}\mathrm{Tr}K(t). \tag{26}$$

A diffusion-based effective dimension is

$$d_{\text{eff}}^{(\text{diff})}(t) = -2\frac{\mathrm{d}\log P_{\text{return}}(t)}{\mathrm{d}\log t}. \tag{27}$$

Independently, if $\Lambda_{\text{shell}}(r)$ is a shell-averaged scalar load around a source, then in a Gaussian scaling window

$$\Lambda_{\text{shell}}(r) \sim \begin{cases} r^{2-d_{\text{eff}}}, & d_{\text{eff}} \neq 2, \\ \log r, & d_{\text{eff}} = 2. \end{cases} \tag{28}$$

This gives a shell-based estimator

$$d_{\text{eff}}^{(\Lambda)}(r) = 2 - \frac{\mathrm{d}\log|\Lambda_{\text{shell}}(r)|}{\mathrm{d}\log r}. \tag{29}$$

The strongest scalar closure condition is therefore

$$d_{\text{eff}}^{(\text{diff})} \simeq d_{\text{eff}}^{(\Lambda)} \simeq 3 \tag{30}$$

in the same scaling window where the timing law fits best. The value 3 is special because a scalar Green profile in three effective dimensions gives the familiar $1/r$ envelope.

4.3 A severe scalar benchmark

A representative severe scalar benchmark is shown in Table 1. It compares the diffusion-based and shell-based dimension estimates in the same candidate window.

Table 1: A scalar closure benchmark. The diffusion estimator can pass while the scalar-shell estimator fails. That outcome does not refute the existence of an internal scalar load, but it does block the stronger claim that the scalar sector has closed as an emergent three-dimensional geometry.

Quantity	Target	Measured value	Status
$d_{\mathrm{eff}}^{(\mathrm{diff})}$	near 3	2.96	pass individually
$d_{\mathrm{eff}}^{(\Lambda)}$	near 3	1.97	fail individually
Joint closure	equal in same window	$2.96 \neq 1.97$	fail
Timing interpretation	same window	not sufficient alone	open

This kind of failure is scientifically valuable. It says that scalar timing can be real as an internal observable while still failing to become geometry. It also prevents the scalar sector from being mistaken for full gravity. A larger multi-substrate scalar benchmark should be included only when the rows are reproduced from a fixed script and a fixed extraction map. The paper should not invent scalar evidence to motivate the tensor test.

5 Why scalar timing is holonomy-free

The scalar timing law assigns a scalar load Λ_i to each vertex. The discrete difference along an oriented edge $i \to j$ is

$$(\mathrm{d}\Lambda)_{ij} = \Lambda_j - \Lambda_i. \tag{31}$$

For a closed loop $\gamma = (i_0, i_1, \ldots, i_m = i_0)$,

$$\sum_{a=0}^{m-1} (\mathrm{d}\Lambda)_{i_a i_{a+1}} = \sum_{a=0}^{m-1} (\Lambda_{i_{a+1}} - \Lambda_{i_a}) = 0. \tag{32}$$

This is the discrete version of the identity $\mathrm{d}^2 = 0$: exact forms are closed. A pure scalar gradient cannot produce nonzero loop holonomy.

Gravity, however, contains more than scalar clock lag. In weak-field language, the metric perturbation contains components that behave like scalar, vector, and tensor sectors under spatial rotations. The radiative content is tensorial. A scalar law may match one clock-readable slice of weak-field behavior, but it cannot generate the transverse-traceless sector by itself.

This is why the spin-2 test is necessary. If QATNU is only a scalar clock model, it is not an emergent-gravity theory. If QATNU produces a tensor sector from its own bond dynamics, it becomes something more.

6 The textbook ladder and the projection trap

Linearized general relativity can be decomposed into scalar, vector, and tensor parts. In transverse-traceless gauge, gravitational waves satisfy

$$\Box h_{ij}^{\mathrm{TT}} = 0, \qquad \partial^i h_{ij}^{\mathrm{TT}} = 0, \qquad \delta^{ij} h_{ij}^{\mathrm{TT}} = 0. \tag{33}$$

The two independent components of h_{ij}^{TT} are the two helicity-2 polarizations.

A substrate model can fall into a projection trap. One defines a tensor observable, applies a transverse-traceless projector, and then observes that the result behaves like a transverse-traceless tensor because the projector enforced that structure. This is not emergence. It is reconstruction.

The TT projector in momentum space is

$$\Pi_{ij}(k) = \delta_{ij} - \frac{k_i k_j}{k^2}, \tag{34}$$

$$P_{ij,kl}^{\mathrm{TT}}(k) = \frac{1}{2}\left[\Pi_{ik}\Pi_{jl} + \Pi_{il}\Pi_{jk} - \Pi_{ij}\Pi_{kl}\right]. \tag{35}$$

This operator is useful as a diagnostic. It can measure how much of a substrate fluctuation lies in the TT sector. But a successful derivation cannot depend on P^{TT} as the reason the field is spin-2. The substrate must generate modes whose low-energy dynamics already select the TT sector.

7 The Pauli-Fierz target

The weak-field kinetic benchmark for a free massless spin-2 field is the Pauli-Fierz structure. For a symmetric tensor $h_{\mu\nu}$ with trace $h = \eta^{\mu\nu} h_{\mu\nu}$, the quadratic Lagrangian density may be written as

$$\mathcal{L}_{\mathrm{PF}} = \frac{1}{2}\partial_\lambda h_{\mu\nu}\partial^\lambda h^{\mu\nu} - \partial_\mu h^{\mu\nu}\partial^\lambda h_{\lambda\nu} + \partial_\mu h^{\mu\nu}\partial_\nu h - \frac{1}{2}\partial_\lambda h \partial^\lambda h. \tag{36}$$

It is invariant under the linearized gauge transformation

$$h_{\mu\nu} \mapsto h_{\mu\nu} + \partial_\mu \xi_\nu + \partial_\nu \xi_\mu. \tag{37}$$

Coupling to a conserved source has the form

$$\mathcal{L}_{\mathrm{int}} = \frac{\kappa_g}{2} h_{\mu\nu} T^{\mu\nu}, \qquad \partial_\mu T^{\mu\nu} = 0. \tag{38}$$

For QATNU, this gives a clear target. The substrate does not need to begin with the Pauli-Fierz Lagrangian. It must produce it, or its lattice/hydrodynamic equivalent, as the long-wavelength effective action of its own fluctuations.

Table 2: Autonomous spin-2 emergence criteria.

Criterion	Required behavior	Failure mode
Mode count	Two gapless low-k tensor modes	Extra scalar/vector modes remain dynamical, or no gapless tensor modes appear
Dispersion	$\Omega^2 = c_g^2 k^2 + O(k^4)$	Massive gap, diffusive mode, unstable mode, or non-wave-like scaling
Transversality	Nonphysical polarizations decouple in the infrared	TT behavior appears only after projection by hand
Propagator	Tensor two-point function scales as P^{TT}/k^2	Flat, short-range, or wrong-power spectrum
Kinetic structure	Effective quadratic action matches Pauli-Fierz ratios up to field normalization	Arbitrary symmetric tensor dynamics with wrong constraint structure
Source coupling	Coupling to a conserved effective source is universal at leading order	Coupling depends on probe species or chosen readout map
Output number	At least one dimensionless number is predicted, not fitted	All quantities are chosen to match known gravity

8 No-go constraints and the required escape hatch

A massless spin-2 sector is not just another collective excitation. The Weinberg-Witten theorem states, in one standard reading, that an ordinary Lorentz-invariant theory with a Lorentz-covariant conserved stress-energy tensor cannot contain composite massless particles of spin greater than one that themselves carry that stress-energy in the usual way. The theorem is not a dismissal of emergent gravity, but it does force the route to be explicit.

For QATNU, the possible escape is not optional decoration. It is load-bearing. The microscopic model is not a continuum Lorentz-covariant quantum field theory on a fixed spacetime. It begins from a finite graph, local update rules, and bond registers. If a Lorentzian spin-2 sector appears, Lorentz symmetry and gauge redundancy must be emergent infrared facts rather than microscopic assumptions. Similarly, the conserved source that couples to the emergent tensor cannot be an ordinary fundamental stress tensor inserted at the microscopic level; it must arise as an effective conserved source of the long-wavelength sector.

This gives two extra tests:

1. **Emergent Lorentz test.** The low-momentum tensor dispersion must approach a single limiting speed, and angular anisotropy must decrease with system size or coarse-graining scale.
2. **Emergent gauge/source test.** Nonphysical polarizations must decouple because of an infrared redundancy or constraint, and the long-wavelength source must be conserved in the same effective limit.

A QATNU spin-2 result that ignores these points would not solve the problem. It would merely move the inconsistency into the assumptions. A successful result must show where the hypotheses of the no-go theorem fail and how the effective theory nevertheless recovers the observed weak-field structure.

9 Where a tensor could come from in QATNU

The scalar load Λ_i is built from the isotropic part of nearby bond occupation. A tensor candidate must instead come from anisotropy, correlations, and directional structure in the bond network. Suppose an edge $e = \langle i, j\rangle$ near coarse-grained position x has a unit tangent vector $\hat{t}_e$ in an embedding or emergent local chart, and let q_e be a bond observable such as centered occupation,

$$q_e = \langle n_e\rangle - \langle n\rangle_{\text{local}}. \tag{39}$$

A natural symmetric traceless internal tensor is

$$M_{ab}(x) = \sum_{e\in\mathcal{N}(x)} q_e \left(\hat{t}_{e,a}\hat{t}_{e,b} - \frac{1}{d}\delta_{ab}\right), \tag{40}$$

where d is the local effective dimension and $\mathcal{N}(x)$ is a coarse-graining neighborhood. This M_{ab} is not yet gravity. It is a candidate substrate anisotropy variable.

The emergence question is whether the connected two-point function of M_{ab} develops the infrared form

$$\langle M_{ab}(k)M_{cd}(-k)\rangle_c \sim \frac{Z}{k^2}P^{\text{TT}}_{ab,cd}(k) + \text{nonuniversal terms}, \tag{41}$$

without building P^{TT} into the definition of M_{ab}. Here Z is a wavefunction normalization. This expression is a diagnostic of emergent spin-2 behavior, not a definition of it.

9.1 The local chart problem

Equation (40) uses edge directions $\hat{t}_e$. On an embedded square or cubic lattice this is straightforward: $\hat{t}_e$ comes from the chosen lattice embedding. But on a general graph, edge directions are not fundamental data. If the chart is chosen arbitrarily, the tensor observable may smuggle in the very geometry the calculation claims to derive.

The tensor protocol must therefore state how local direction data are obtained. There are three acceptable levels, each with a different claim strength.

1. **Embedded-lattice validation.** For square, triangular, cubic, or other explicitly embedded patches, $\hat{t}_e$ is inherited from the known embedding. These tests validate code, polarization analysis, and finite-size scaling, but they do not by themselves prove background-free geometry.
2. **Spectral or diffusion-map charts.** For non-embedded graphs, local coordinates can be reconstructed from low graph-Laplacian eigenmodes, diffusion distances, or another predeclared intrinsic charting algorithm. In this case $\hat{t}_e$ is an output of the graph operator, not an external drawing.
3. **Chart-robust intrinsic test.** A candidate spin-2 signal must remain stable under reasonable chart reconstructions. If the apparent TT sector changes drastically when the charting method changes, the result is chart artifact, not tensor emergence.

The failure rule is simple:

> If the apparent TT sector depends sensitively on the arbitrary embedding or local chart used to define $\hat{t}_e$, then the tensor sector has been inserted through the coordinate choice rather than generated by QATNU.

9.2 Hydrodynamic logic

A gapless long-wavelength mode usually requires a conservation law, a broken symmetry, a gauge redundancy, a constraint, or a critical point. QATNU must identify the mechanism that protects a massless spin-2 sector. Plausible routes include constraint-protected bond stress, critical tensor correlations, entanglement-elasticity modes, and gauge redundancy from relational encoding. These are mechanisms to test, not conclusions to assume.

Table 3: Candidate protection mechanisms for a massless tensor sector. The point is to prevent "gapless spin-2" from becoming a label attached after the fact. Each route predicts a different numerical failure pattern.

Candidate mechanism	Numerical signature	Falsifier
Constraint-protected bond stress	Approximate Ward identity; nonphysical polarizations suppressed as $k \to 0$	Scalar/vector leakage remains finite in the infrared
Critical tensor correlations	Correlation length grows near a promotion transition; tensor correlator shows finite-size scaling	Apparent $1/k^2$ tail vanishes away from a tuned finite-size point
Entanglement elasticity	Shear-like anisotropic modes propagate with stable c_g	Modes are diffusive, damped, or ordinary lattice phonons only
Relational gauge redundancy	Multiple microscopic bond descriptions map to the same infrared tensor state	Tensor field depends on arbitrary encoding choices or gauge-fixing conventions

10 Criticality rather than plausibility

The tensor question should not be phrased as "can one define a tensor observable?" That is too weak. Many systems allow tensor-valued bookkeeping. The real question is whether a microscopic tensor disturbance dies out, reproduces itself across scales, or runs away.

Define a coarse-graining map $\mathcal{C}_b$ that sends microscopic tensor data to a field at scale b. For a fixed microscopic injection of bond anisotropy, define the tensor reproduction factor

$$\mathcal{R}_{\mathrm{TT}}(b) = \frac{\text{TT tensor power after coarse-graining by scale } b}{\text{TT tensor power injected at microscopic scale}}. \tag{42}$$

More explicitly, if $S_{ab,cd}(k)$ is the raw tensor correlator and P^{TT} is used only diagnostically, one

may use

$$\mathcal{R}_{\mathrm{TT}}(b;k_0)=\frac{\int_{|k|<k_0/b}\mathrm{Tr}\left[P^{\mathrm{TT}}(k)S_b(k)\right]\,\mathrm{d}^dk}{\int_{|k|<k_0}\mathrm{Tr}\left[P^{\mathrm{TT}}(k)S_{\mathrm{inj}}(k)\right]\,\mathrm{d}^dk}, \tag{43}$$

where S_b is measured after coarse-graining and S_{inj} describes the predeclared injected perturbation or initial anisotropy.

The regimes are:

$$\mathcal{R}_{\mathrm{TT}}<1 \quad \text{subcritical: tensor structure dies out; no emergent graviton,} \tag{44}$$

$$\mathcal{R}_{\mathrm{TT}}\simeq 1 \quad \text{critical: tensor structure survives scale change; possible spin-2 sector,} \tag{45}$$

$$\mathcal{R}_{\mathrm{TT}}>1 \quad \text{supercritical/runaway: tensor growth is unstable, not weak-field gravity.} \tag{46}$$

This reframes the spin-2 search as a critical phenomenon. Scalar timing may be real while remaining subcritical as gravity. Tensor anisotropy may be definable while failing to reproduce itself across scales. The discovery question is whether QATNU reaches an infrared regime in which tensor anisotropy becomes self-sustaining, gapless, and constrained rather than merely present.

11 The numerical spin-2 protocol

The first severe test can be numerical. It should begin with the finite QATNU Hamiltonian, then expand system size and topology while measuring tensor correlations. The protocol has to be severe enough to fail.

11.1 Pre-registration before simulation

Before inspecting tensor spectra, the following choices must be fixed. Otherwise the calculation can drift into post-selection: try many observables, windows, charts, and graph families until something looks like a graviton.

Table 4: Pre-registered choices for a Mercury-style spin-2 run. Changes after seeing spectra should be reported as exploratory, not confirmatory.

Choice	Must be fixed before the run
Hamiltonian	H, $\chi_{\max}$, bond-energy ladder ϵ_m, ω, λ scan, boundary conditions
Graph family	path/cycle controls, embedded patches, random regular controls, periodic lattices, system sizes
Raw observables	Λ_i, M_{ab}, $S_{ab,cd}(k,t)$, source perturbations, clock observables
Chart/coarse graining	embedding or intrinsic chart algorithm, neighborhood $\mathcal{N}(x)$, scale factor b
Fit windows	k window, Ω window, time window, finite-size sequence
Success thresholds	allowed $m_g(L)$ drift, TT fraction, anisotropy residual, Ward residual, R_g stability
Null controls	which failed substrates must fail before candidate lattices are interpreted

Changing these choices is allowed in exploratory work, but then the output is hypothesis generation, not a Mercury-style result. A serious claim must separate exploratory tuning from confirmatory runs.

11.2 Step 1: choose graph families

Use graph families that separate obvious finite-size artifacts from robust scaling behavior:

- paths and cycles for baseline one-dimensional behavior;
- square, triangular, and cubic patches for embedded lattice comparison;
- random regular graphs for bounded-degree non-geometric controls;
- star, bowtie, diamond, and bottleneck graphs as stress tests;
- increasingly large periodic lattices for low-momentum scaling.

A result that appears only on a tiny path graph is not evidence of emergent gravity. The relevant question is whether the same infrared structure strengthens with size and survives moderate topology changes.

11.3 Step 2: define substrate observables before projection

Before applying any spin decomposition, define raw observables:

$$\Lambda_i = \sum_{e \ni i} \log(1 + \langle n_e \rangle), \tag{47}$$

$$M_{ab}(x) = \sum_{e \in \mathcal{N}(x)} q_e \left(\hat{t}_{e,a} \hat{t}_{e,b} - \frac{1}{d} \delta_{ab} \right), \tag{48}$$

$$S_{ab,cd}(k,t) = \langle M_{ab}(k,t) M_{cd}(-k,0) \rangle_c. \tag{49}$$

The tensor candidate must be defined by substrate quantities alone: bond occupations, bond correlations, update rules, local direction data, and coarse-graining neighborhoods. The TT projection is reserved for analysis.

11.4 Step 3: measure spectra

Compute the dynamical structure factor

$$S_{ab,cd}(k,\Omega) = \int \mathrm{d}t \, e^{i\Omega t} S_{ab,cd}(k,t). \tag{50}$$

The desired signature is a pair of sharp low-k peaks with

$$\Omega(k) = c_g |k| + O(k^3) \tag{51}$$

whose polarization content is dominantly TT.

A diffusive scalar mode would instead show

$$\Omega \sim -iDk^2. \tag{52}$$

An ordinary gapped excitation would show

$$\Omega^2 = m^2 + c^2 k^2 + \cdots . \tag{53}$$

Neither counts as massless spin-2 gravity.

11.5 Step 4: fit the infrared propagator

The equal-time tensor correlator should be fit to

$$S_{ab,cd}(k) \simeq \frac{Z}{k^2 + m_g^2} P^{\mathrm{TT}}_{ab,cd}(k) + R_{ab,cd}(k), \tag{54}$$

where R contains analytic short-range and non-TT contributions. A successful massless tensor sector requires

$$m_g \to 0 \tag{55}$$

with increasing system size, while Z remains finite and the TT fraction increases in the infrared.

11.6 Step 5: test source coupling

Introduce controlled source perturbations. A source could be a localized change in matter frustration, a controlled bond-bias pattern, or a localized graph defect. Let $S^{\mu\nu}$ denote the effective conserved source inferred from the microscopic perturbation. The candidate tensor response should obey a lattice analogue of

$$\mathcal{D}^{\mu\nu\alpha\beta} h_{\alpha\beta} = \kappa_g S^{\mu\nu}, \tag{56}$$

where $\mathcal{D}$ reduces to the Pauli-Fierz kinetic operator at low momentum. Universality means that different microscopic probes that share the same conserved long-wavelength source produce the same leading tensor response.

11.7 Step 6: extract the primary Mercury number

At least one number must be predicted by the substrate rather than fixed by matching. The primary target should be chosen before the run. For QATNU, the cleanest first candidate is

$$R_g = \frac{c_g}{c_{\mathrm{clock}}}, \tag{57}$$

where c_g is the infrared tensor-wave speed extracted from the spin-2 structure factor, and c_{clock} is the propagation speed or causal-cone speed of the scalar clock sector measured in the same microscopic units.

This is the natural Mercury number because it is dimensionless, internal to the substrate, and not fixed by external unit choice. If R_g is forced to equal one by construction, it is not an output. If it approaches one only after independent extraction of both speeds, that is a nontrivial result. If it approaches a different stable value, the model has made a real prediction and must explain the resulting Lorentz mismatch.

Secondary outputs should still be recorded, but they should not replace the primary target:

- the infrared tensor stiffness Z in units of the microscopic bond scale;
- the ratio between scalar clock susceptibility and tensor coupling, α/κ_g;
- a critical exponent for the tensor-correlation length near a promotion transition;
- a post-Newtonian parameter such as γ_{PPN}, if the substrate reaches a static weak-field limit.

A model that chooses these values to match general relativity has not passed the Mercury criterion. A model that outputs them and then survives comparison has.

11.8 Finite-size and infrared diagnostics

The protocol should record diagnostic quantities, not only plots. Let L be a linear system size or effective diameter. A minimal diagnostic table should include:

$$m_g(L) \quad \text{effective tensor mass gap,} \tag{58}$$

$$F_{\mathrm{TT}}(k,L) = \frac{\mathrm{Tr}[P^{\mathrm{TT}}(k)S(k,L)]}{\mathrm{Tr}[S(k,L)]}, \tag{59}$$

$$\Delta_c(k,L) = \frac{\mathrm{std}_{\hat{k}}[c_g(\hat{k},L)]}{\mathrm{mean}_{\hat{k}}[c_g(\hat{k},L)]}, \tag{60}$$

$$W(k,L) = \frac{\|k^a S_{ab,cd}(k,L)\|}{\|S_{ab,cd}(k,L)\|}, \tag{61}$$

$$\mathcal{R}_{\mathrm{TT}}(b,L) \quad \text{coarse-graining reproduction factor.} \tag{62}$$

Here F_{TT} is the TT fraction, Δ_c measures angular anisotropy in the emergent tensor speed, and W is a Ward-residual or transversality-failure diagnostic.

The desired finite-size trends are

$$m_g(L) \downarrow 0, \qquad F_{\mathrm{TT}}(k,L) \uparrow 1, \qquad \Delta_c(k,L) \downarrow 0, \qquad W(k,L) \downarrow 0, \qquad \mathcal{R}_{\mathrm{TT}} \to 1. \tag{63}$$

A one-size result is not enough. The signal must improve, not degrade, when the system is made more capable of supporting an infrared sector.

12 Phase-0 null controls and minimal failed substrates

The protocol must be able to say no. Several graph families fail before any expensive computation is needed. They are not embarrassments; they are calibration cases.

Table 5: Minimal failure controls. These cases are useful precisely because they should not pass the spin-2 test. A protocol that cannot reject them is not severe enough to confirm a harder substrate.

Substrate family	Why it fails the spin-2 target	What it can still test
One-dimensional path	No transverse tensor sector exists; loop holonomy and TT polarization are absent	Scalar timing law, finite-size clock shifts, Lieb-Robinson cone
One-dimensional cycle	Periodic boundary adds a loop but not transverse-traceless helicity-2 content	Scalar closure and topological sensitivity
Star graph	No smooth low-k continuum or translation-invariant infrared sector; hub dominates spectrum	Sensitivity to bottlenecks and source placement
Pure scalar observable model	Λ_i can fit clock shifts but cannot produce helicity-2 modes	Negative control for the projection trap
Hand-projected tensor control	TT content appears because the projector imposed it	Positive diagnostic control, not evidence of emergence

These controls should be run first. A minimal Phase-0 report should list F_{TT}, m_g trend, $\mathcal{R}_{\mathrm{TT}}$, and verdict for each control. Expected outcomes are shown in Table 6.

Table 6: Phase-0 null controls to be run before interpreting candidate spin-2 graphs. Values are not filled here because the paper should not report results that have not been generated; the point is to pre-register the verdict logic.

Graph/control	Expected F_{TT}	Expected m_g trend	Expected $\mathcal{R}_{\mathrm{TT}}$	Verdict
Path	undefined or zero	no tensor sector	< 1	fail
Cycle	undefined or zero	no tensor sector	< 1	fail
Star	no low-k continuum	hub dominated	unstable or meaningless	fail
Hand-projected control	high by construction	diagnostic only	invalid	projection trap

A small path or star graph can still be useful for debugging the Hamiltonian, measuring scalar clock shifts, and validating code. It should not be described as evidence for emergent gravity. The first nontrivial spin-2 search requires a graph family with an approximate continuum limit, enough local angular structure to define anisotropic tensor variables, and finite-size scaling that improves rather than degrades as the system grows.

13 Analytical route: fast-edge elimination

A useful analytical route is to integrate out fast coupler or bond variables and see what effective operator acts on slow observables. Suppose slow variables s and fast variables f obey a linearized system

$$\dot{s} = A_{ss}s + A_{sf}f + \eta_s, \tag{64}$$

$$\dot{f} = A_{fs}s + A_{ff}f + \eta_f, \tag{65}$$

with $\|A_{ff}\|$ setting the fast timescale. Markovian adiabatic elimination gives

$$f \simeq -A_{ff}^{-1}A_{fs}s, \tag{66}$$

so

$$\dot{s} \simeq (A_{ss} - A_{sf}A_{ff}^{-1}A_{fs})s + \tilde{\eta}_s. \tag{67}$$

The effective slow generator is the Schur complement of the fast block.

For a scalar slow variable, this elimination can produce the graph-Laplacian closure. For a tensor slow variable, the question is whether the same kind of elimination produces a Pauli-Fierz-like operator rather than merely another scalar diffusion equation.

A general low-momentum inverse propagator can be expanded schematically as

$$D^{\text{eff}}_{ab,cd}(\Omega, k) = A_{0,ab,cd} + A^{ij}_{2,ab,cd}k_ik_j + A_{\Omega,ab,cd}\Omega^2 + \cdots . \tag{68}$$

In a locally isotropic sector it can be decomposed into scalar, vector, tensor, and TT projectors:

$$D_{\text{eff}} = D_S P_S + D_V P_V + D_T P_T + D_{\text{TT}} P_{\text{TT}} + \text{mixing terms}. \tag{69}$$

The pass condition is not merely that some tensor-looking term appears. The required infrared behavior is

$$D_{\text{TT}}(\Omega, k) \sim Z^{-1}(\Omega^2 - c_g^2 k^2), \tag{70}$$

while nonphysical sectors are constrained, gapped, or decouple. The analytical fail condition is

$$D_{\text{eff}} \sim \kappa L_w + \mu^2 I \tag{71}$$

in every slow channel. In that case the substrate produces scalar diffusion/screening, not spin-2 gravity.

If the Schur complement produces only screened scalar operators of the form $\kappa L_w + \mu^2 I$, the spin-2 program fails at the analytical level. If it produces the TT inverse propagator above in the infrared, the numerical test becomes urgent.

14 Internal fields versus external signals

Even if QATNU produces an internal tensor mode, external detectability is a separate question. For scalar loads, the readout map is $\Lambda_{\text{ext}}(x) = \sum_i C(x, i)\Lambda_i$. For a tensor candidate, an external readout would require a tensor coupling kernel,

$$h^{\text{ext}}_{ab}(x) = \sum_y C_{ab}{}^{cd}(x, y)M_{cd}(y). \tag{72}$$

Possible kernel classes include:

1. **Null kernel.** The internal field has no measurable external response.
2. **Near-field kernel.** The signal decays exponentially with distance from the substrate.
3. **Substrate-mediated kernel.** The signal is carried by ordinary mechanical, thermal, dielectric, or electromagnetic channels.
4. **Boundary-modulated kernel.** The internal state modulates cavity or impedance conditions.
5. **Vacuum-like Green kernel.** The response has a power-law tail not reducible to mundane substrate coupling.

The strongest claim requires the last class, but the first four are more likely experimentally. A null or mundane external kernel would not invalidate internal spin-2 emergence as a mathematical phenomenon, but it would limit the claim as physics.

This paper does not need to solve the external-kernel problem in order to define the internal spin-2 test. It does need to keep the distinction explicit. A substrate-generated tensor mode is not yet a detected gravitational field until a coupling channel is specified or bounded.

15 What the framework solves

The framework does not solve quantum gravity by assertion. Its value is more specific.

15.1 It separates projection from derivation

The scalar timing law is treated as a projection: a clock-readable internal observable. It is not inflated into a full gravitational theory. The spin-2 sector is reserved for a separate derivation.

15.2 It gives the model a hard target

Rather than adding more interpretive layers, the framework defines a narrow pass/fail target: autonomous helicity-2 emergence. This prevents endless movement of the goalposts.

15.3 It makes failure informative

If scalar closure fails, the scalar sector narrows. If tensor modes fail, QATNU remains an interesting internal-medium model but not emergent gravity. If external kernels are null, the field remains internal. Each failure reduces exaggeration and increases clarity.

15.4 It identifies the first possible discovery point

The first possible discovery is not a new metaphor for redshift. It is a substrate-generated spin-2 sector with a nontrivial output number. That is the point at which the model would begin to compete with existing physics rather than redescribe it.

16 What would count as success

A successful QATNU spin-2 derivation would produce the following chain:

$$H_{\text{micro}} \longrightarrow \text{bond correlations} \longrightarrow M_{ab} \longrightarrow h_{ab}^{\text{TT}} \longrightarrow \mathcal{L}_{\text{PF}} + h_{\mu\nu}T^{\mu\nu} + \text{output number}. \tag{73}$$

Each arrow must be earned.

- $H_{\text{micro}} \to$ bond correlations: follows from unitary QATNU evolution.
- bond correlations $\to M_{ab}$: follows from a substrate-defined tensor observable.
- $M_{ab} \to h_{ab}^{\text{TT}}$: must arise in the infrared, not by definition.
- $h_{ab}^{\text{TT}} \to \mathcal{L}_{\text{PF}}$: requires Pauli-Fierz kinetic ratios and gauge-like decoupling.
- coupling to $T^{\mu\nu}$: requires universal long-wavelength source response.
- output number: requires a prediction not chosen to match GR, with $R_g = c_g/c_{\text{clock}}$ as the primary first target.

A concise success statement would be:

> In the low-energy, long-wavelength limit of QATNU, the connected anisotropic bond-stress correlator develops two gapless transverse-traceless modes with Pauli-Fierz dynamics and universal conserved-source coupling, and predicts a dimensionless infrared number not fixed by the matching procedure.

17 What would count as failure

The program should also state failure conditions clearly.

1. **Scalar-only closure.** The model supports Λ_i and the timing law, but no independent tensor sector.
2. **Projection-only tensor.** A TT field appears only after applying the projector to a field that does not itself propagate as spin-2.
3. **Wrong spectrum.** The tensor candidate is gapped, diffusive, unstable, or dominated by scalar/vector modes.
4. **Wrong propagator.** The correlator lacks a $1/k^2$ infrared tensor pole.
5. **Wrong coupling.** Different probes couple differently at leading order, preventing universal gravitational interpretation.
6. **No output number.** The model can match known formulas only by choosing parameters after the fact.
7. **No no-go escape.** The model claims a composite graviton while retaining the assumptions that forbid it, or fails to show emergent Lorentz/gauge structure in the infrared.
8. **Null externalization.** Internal modes exist but cannot couple to any external measurement channel except ordinary material effects.

9. **Chart dependence.** The apparent TT sector changes qualitatively with the chosen embedding or intrinsic chart.
10. **Post-selected protocol.** The result appears only after changing observables, fitting windows, or thresholds in response to the spectra.
11. **Criticality failure.** $\mathcal{R}_{\mathrm{TT}}$ remains subcritical or becomes runaway rather than approaching a stable critical value.

These are not rhetorical losses. They are the conditions that make the proposal scientific.

18 Discussion

The strongest version of the framework is not the claim that scalar timing is gravity. It is the claim that scalar timing is a disciplined first observable, while gravity requires a further tensor derivation. This distinction matters because many speculative programs become unfalsifiable by treating every partial resemblance as confirmation. A clock lag resembles gravitational redshift; a graph Laplacian resembles spatial geometry; a projected tensor resembles linearized gravity. Each resemblance is meaningful, but none is decisive by itself.

The QATNU proposal becomes serious only where resemblance ends and derivation begins. Its microscopic Hamiltonian supplies a lawful substrate. Its bond registers supply internal structure. Its scalar load supplies a clock-readable sector. But the gravitational question is whether anisotropic bond correlations can organize themselves into a protected massless spin-2 field while also explaining why the usual no-go assumptions do not apply in the emergent limit.

The criticality layer is important because it prevents a weaker question from replacing the real one. Tensor patterns may exist locally and still be irrelevant in the infrared. A graviton-like sector would require scale survival: not merely tensor definition, but tensor reproduction under coarse-graining, with nonphysical polarizations constrained and the tensor mass gap falling away.

The answer may be no. If the answer is no, the model still teaches something: a scalar timing sector can be mathematically coherent and experimentally discussable without being gravity. If the answer is yes, the model enters a different category: it would provide a concrete route from local quantum automaton dynamics to the weak-field tensor content of general relativity.

The Mercury rung is therefore not a slogan. It is a stopping rule. Either the substrate generates the spin-2 sector and outputs a number, or the gravitational claim stops there.

19 Conclusion

This paper has presented a QATNU-based emergent-gravity test. The proposal begins with a finite local quantum substrate: matter qubits on vertices, bond registers on edges, and frustration-controlled bond promotion. It defines an internal scalar load, a clock law, graph-Laplacian scalar closure, and effective dimension tests. It then states why these scalar structures are insufficient for gravity: scalar timing is holonomy-free and cannot produce the transverse-traceless helicity-2 sector of weak-field general relativity.

The decisive target is autonomous spin-2 emergence. QATNU must produce two gapless TT modes,

a Pauli-Fierz-like kinetic operator, a $1/k^2$ tensor propagator, universal conserved-source coupling, a clear escape from the Weinberg-Witten obstruction, and at least one nontrivial output number. The primary first number is $R_g = c_g/c_{\text{clock}}$. The tensor sector must be diagnosed by projection, not created by projection. That is the difference between a derivation and a reconstruction.

The hardened protocol adds three constraints that are essential rather than decorative: pre-registration before simulation, chart robustness for the tensor observable, and criticality under coarse-graining through $\mathcal{R}_{\text{TT}}$. These additions make the test less forgiving and therefore more useful.

If QATNU passes this test, it becomes a genuine candidate for emergent gravitational physics. If it fails, the scalar timing sector remains a useful controlled projection or internal-medium theory, but not a derivation of gravity. Either outcome is progress because the target is now sharp enough to say no.

A Notation registry

Table 7: Notation registry.

Symbol	Meaning	Level
H_{micro}	Microscopic Hamiltonian	Vertex and bond substrate
$G = (V, E)$	Graph of sites and bonds	Discrete substrate
$W = (w_{ij})$	Weighted adjacency matrix	Internal effective graph
$L_w = D - W$	Weighted graph Laplacian	Scalar closure and diffusion
ρ	Effective source density	Internal source
Λ_i	Scalar load at vertex i	Internal observable
$\Lambda_{\text{ext}}(x)$	Externally read scalar response	Sensor-level observable
$C(x, i)$	Internal-to-external scalar kernel	Coupling map
ω	Bare clock frequency	Microscopic/observable baseline
ω_{eff}	Effective clock frequency	Timing observable
α	Clock susceptibility	Effective parameter
M_{ab}	Bond-anisotropy tensor candidate	Internal tensor observable
$\hat{t}_e$	Local edge direction	Embedded or reconstructed chart data
$h_{\mu\nu}$	Metric perturbation benchmark	Continuum weak-field gravity
h_{ij}^{TT}	Transverse-traceless tensor mode	Radiative spin-2 sector
c_g	Emergent tensor-wave speed	Output candidate
c_{clock}	Clock/causal propagation speed	Internal scalar/clock sector
R_g	c_g/c_{clock}	Primary Mercury number
Z	Tensor correlator normalization	Output candidate

Symbol	Meaning	Level
κ_g	Effective gravitational coupling	Output/matching parameter
$\mathcal{R}_{\mathrm{TT}}$	TT reproduction factor under coarse-graining	Criticality diagnostic
F_{TT}	TT fraction of tensor correlator	Diagnostic
$W(k, L)$	Ward/transversality residual	Diagnostic

B Minimal implementation sketch

A minimal exact-diagonalization loop is:

1. construct the graph and edge list;
2. build the Hamiltonian from the matter, bond, and interaction terms;
3. diagonalize or time-evolve selected states;
4. compute $\langle n_e(t)\rangle$, $\Lambda_i(t)$, and $M_{ab}(x,t)$;
5. compute scalar closure diagnostics;
6. compute tensor structure factors, TT fractions, Ward residuals, and chart robustness diagnostics;
7. fit low-k dispersion and propagators;
8. extract c_g, c_{clock}, and R_g independently;
9. run Phase-0 null controls before interpreting candidate lattices;
10. repeat over system size, topology, λ, and predeclared chart/coarse-graining rules.

A confirmatory run should generate a machine-readable record of the pre-registration table, code version, random seeds, graph definitions, fit windows, and all failure thresholds. This is not bureaucratic overhead. It is how the Mercury criterion prevents a diagnostic projector or a chosen fit window from becoming the hidden source of the result.

C Scalar closure benchmark appendix

The scalar sector is not being abandoned because it is meaningless. It is being bounded because its strongest closure test has not yet passed. The benchmark in Table 1 shows the relevant pattern: one dimension estimator may sit near 3 while another independent estimator does not. That is exactly the kind of failure the program should report rather than conceal.

A fuller scalar benchmark table should include one row per substrate family or extraction map, with at least the following fields:

$$\Big(G,\ W\text{-map},\ d_{\mathrm{eff}}^{(\mathrm{diff})},\ d_{\mathrm{eff}}^{(\Lambda)},\ \text{timing residual},\ \text{window overlap},\ \text{verdict}\Big)\,. \tag{74}$$

If all rows fail, that is not a reason to inflate the scalar claim. It is a reason to keep scalar timing as an internal observable and move the gravitational claim to the spin-2 test.

D References

1. A. Einstein, "Die Feldgleichungen der Gravitation," *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, 1915.
2. M. Fierz and W. Pauli, "On relativistic wave equations for particles of arbitrary spin in an electromagnetic field," *Proceedings of the Royal Society A*, 1939.
3. S. Weinberg, "Photons and gravitons in perturbation theory: Derivation of Maxwell's and Einstein's equations," *Physical Review*, 1965.
4. S. Weinberg and E. Witten, "Limits on massless particles," *Physics Letters B*, 1980.
5. E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," *Communications in Mathematical Physics*, 1972.
6. C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*. W. H. Freeman, 1973.
7. R. M. Wald, *General Relativity*. University of Chicago Press, 1984.
8. F. Chung and S.-T. Yau, "Discrete Green's functions," *Journal of Combinatorial Theory, Series A*, 2000.
9. T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Physical Review Letters*, 1995.
10. B. Swingle, "Entanglement renormalization and holography," *Physical Review D*, 2012.
11. C. M. Will, *Theory and Experiment in Gravitational Physics*. Cambridge University Press, revised edition.