

How Good Can a Clock Ever Be?

Brecht Corbeel

Absolute Digital Publishers

absolutedigitalpublishers.com

September 20, 2026

Abstract. Quantum mechanics prices a clock’s resolution in energy spread; general relativity charges every clock a self-redshift rent on the energy it contains. The heuristic lineage that combines them — Salecker–Wigner, Ng–van Dam, Lloyd — never performed the accounting that decides what is a limit and what is a calibration: the *mean* self-redshift of a clock’s mass–energy is a systematic, removable in principle to any accuracy, and only the *quantum fluctuations* of the energy that both gravitates and correlates with the tick record dephase the clock irreducibly. Carrying that split through yields a floor on fractional accuracy $\mathcal{A}(R, \tau)$ that selects an optimal clock size through the light-crossing constraint, lands tens of orders of magnitude below today’s 10^{-19} optical clocks (the exponent computed, and bracketed honestly between two bounds on which energy counts), and turns the popular Planck-foam accuracy claims into graded conjectures: we reconstruct the $t_P^{2/3} \tau^{1/3}$ derivations, name their counting assumptions, add the mean-versus-fluctuation objection to the published critiques, and note what interferometric and astrophysical data already bound. The map of every flagship instrument against the floor closes the Quantum–Relativity Papers: ten papers of operational quantum-relativity, ending on how much room the two theories jointly leave — and how little of it we have used.

Keywords: quantum speed limits; clock accuracy; gravitational self-redshift; optical clocks; nuclear clock; spacetime foam bounds; fundamental limits of timekeeping.

Series: The Quantum–Relativity Papers, No. 10.

1 A perfect clock is forbidden twice, and the two prohibitions have never been added up carefully

Ask how good a clock can be and the question sounds like it has an engineering answer: sharpen the oscillator, lengthen the interrogation, shield the environment, and the fractional accuracy $\mathcal{A} = \delta f/f$ keeps falling. The record encourages the reading. A single-ion $^{27}\text{Al}^+$ quantum-logic clock holds a systematic fractional uncertainty of 9.4×10^{-19} [14]; a strontium optical lattice has reached 8.1×10^{-19} [2]; a millimetre-tall column of strontium atoms at JILA resolves a fractional frequency of 7.6×10^{-21} across its own height [11]. Eighteen digits, and the curve has not visibly bent. This paper argues that it must bend eventually, works out the shape of the wall it bends toward, and — as the tenth and closing paper of this series — reads that wall against the Planck-scale claims that have long been attached to it.

The wall exists because two independent parts of physics charge a clock for the same thing, the energy it must contain in order to tick, and the two charges do not cancel. The first is quantum mechanical. Resolving a phase requires an evolution fast enough to distin-

guish states, and the Mandelstam–Tamm and Margolus–Levitin speed limits fix the exchange rate between how fast a system can evolve and the spread of energy it carries [23, 49, 50]: a sharper tick demands a larger energy uncertainty σ_E . We call this the quantum price, and Section 2 assembles it, together with the resource-theoretic accounting of what an autonomous clock must dissipate to run at all [29, 55], as one column of a ledger.

The second charge is relativistic and is the one the literature has handled loosely. Every clock sits in the gravitational potential of its own mass-energy, and its rate is redshifted there. Section 3 makes the charge precise and, in doing so, states this paper’s central technical claim: the redshift from the clock’s *mean* mass-energy is a fixed systematic — large, but a number, and therefore calibrable in principle and no bound on accuracy at all — while the redshift from the *quantum fluctuation* of that mass-energy is not calibrable, because it is correlated with the very tick record the clock is trying to produce. Only the fluctuation is irreducible, and it is bounded below by the quantum price of the previous paragraph: the same σ_E that buys resolution pays the irreducible part of the rent. The cross term that couples the two is ours, and it is exact in the model we state.

That the mean and the fluctuation must be separated is exactly the step the well-known heuristic lineage skips. Salecker and Wigner asked what mass a clock needs so

that its own quantum spread does not blur the interval it measures [79]; Ng and van Dam turned that reasoning, with a light-crossing and a Schwarzschild constraint, into the claim that a length carries an intrinsic uncertainty $\delta l \propto l^{1/3}$ [60, 61]; Lloyd counted a system’s total energy as the resource setting its maximum operation rate [46]; Amelino-Camelia proposed reading such a “foam” off gravitational-wave interferometers [4]. Each of these counts the mean mass-energy as the cost. We report the lineage as the reason the question is worth asking, and we grade its conclusions rather than repeat them: the holography-motivated $\tau^{1/3}$ clock bound has published objections on measurement-scheme grounds [7] and is excluded, in its length form, by astrophysical wavefront-coherence data [72]. Our own objection, developed in Section 6, is the accounting one: the lineage bounds accuracy with a quantity that is calibrable.

This paper therefore makes four claims, each graded where it is made. (a) The systematics-versus-fluctuations accounting of Section 3: the self-redshift from the mean does not bound accuracy, and the floor comes only from the conditionally-fluctuating part of the clock’s own gravitating energy; we derive the cross term (exact in the model, with the choice of which subsystem’s energy enters bracketed both ways in Appendix A). (b) The resulting floor $\mathcal{A}_{\min}(R, \tau)$ and its scaling in Section 4: the rent falls as $1/R$ while a coherent tick caps the size at the light-crossing distance $R \leq c\delta t$, and the two together select an optimal clock size and an absolute floor whose suppression is quadratic in the Planck time, not the one-third power the foam lineage predicts (ours). (c) The map of Section 5: every demonstrated clock, the near-term nuclear and entangled projections, and the floor family on one plane (ours as compilation, every plotted number verified). (d) The graded verdict of Section 6 on the Planck-foam accuracy claims, each grade argued against the published data.

The floor lands far below anything a laboratory has built or projected. Its point is its exponent, not its coefficient: at laboratory sizes it sits scores of orders of magnitude under today’s best clocks, the exact figure quoted from a seeded computation in Section 4 with the order-honesty such a number demands. That distance is the substance of the closing verdict. A gap of that size means the floor is not something a clock is about to run into, and it means the popular reading — that a good enough clock would be reading quantum spacetime — has the accounting backwards: what current clocks approach is the calibrable mean, not the irreducible fluctuation, and the two differ by the whole width of the map.

The series-closer frame is not decoration. The gravitational ingredients here — the operational treatment of self-redshift and of proper time as a measured, not postulated, quantity — are the ones the earlier papers of this series built [20, 21], and the operational stance toward what a measurement of a relativistic quantity actually fixes is the one developed at Corbeel [22]. The

redshift itself remains a contested thing to claim one has measured: whether an atom interferometer tests the gravitational redshift at the Compton frequency [58] or instead tests the universality of free fall was disputed in print [93] and defended [36], and that the dispute turns on *what quantity a reading fixes* is precisely the discipline this paper brings to the clock’s own energy. Ten papers of operational quantum-relativity end, here, on the full landscape of how good a clock can be — and on saying exactly what a clock would have to do to read spacetime rather than calibrate it.

2 Resolution is bought with energy spread, and the speed limits fix the price

A clock is a system whose state runs through a sequence of distinguishable configurations at a known rate; reading it is telling those configurations apart. Both halves of that sentence are quantum-mechanically costly, and the cost is charged in the same currency — energy — which is why the two prices of this paper can be added at all. This section assembles the quantum charge as the first column of the ledger, attributing each piece and marking what the resource-theoretic treatment adds that the bare speed limits do not.

Proposition 1 (The quantum price of a tick; exact where attributed). *A clock reading its oscillation to phase resolution $\delta\phi$ over an interrogation time δt must carry, in the degree of freedom that oscillates, an energy uncertainty bounded below by the Mandelstam–Tamm relation*

$$\sigma_E \geq \frac{\hbar \delta\phi}{2\delta t}, \quad (1)$$

equivalently a quantum-speed-limit time $\tau_{\text{QSL}} = \pi\hbar/(2\sigma_E)$ below which no distinguishable tick can form [23, 49]; the Margolus–Levitin bound $\tau_{\text{QSL}} = \pi\hbar/(2E)$ replaces the variance by the mean energy above the ground state [50], and the attainable rate is the tighter of the two, $\tau_{\text{QSL}} = \max\{\pi\hbar/(2\sigma_E), \pi\hbar/(2E)\}$ [23]. Averaging N independent such degrees of freedom for a time τ lowers the fractional instability as $\sigma(\tau) \propto 1/(\sqrt{N}\tau)$ at the standard quantum limit and as $1/(N\tau)$ at the Heisenberg limit [32]. Any autonomous realization pays, in addition, an entropy production per tick, so that its accuracy is bounded by the entropy it dissipates [29, 55].

The inequality (1) is the oldest piece and the one most often misread. Mandelstam and Tamm derived it as a statement about how fast an expectation value can change: the rate at which any observable evolves is bounded by the product of its own spread and the spread of the Hamiltonian, so a system whose energy is sharp evolves slowly and one that evolves quickly must have

energy that is not sharp [49]. Read as a time-energy uncertainty relation it invites the wrong picture, of a duration one cannot measure precisely; Aharonov and Bohm settled that reading by exhibiting a measurement scheme that beats the naive $\Delta E \Delta t \gtrsim \hbar$ for a measurement’s *duration*, which is what forces the correct statement that δt in Eq. (1) is an evolution time-scale, the time to reach a distinguishable state, not the time one spends looking [3]. Busch’s survey is the authority for the distinction: there is no single time-energy relation with the standing of the position-momentum one, only a family of legitimate statements each valid in its own setting [15], and Eq. (1) is the member that applies to a clock. We take it in exactly that member’s form, and grade its use here exact: it is a theorem about the oscillator, not an analogy.

Margolus and Levitin sharpened the bound in a way that matters for the accounting to come. Their result depends only on the mean energy above the ground state, E , not on its variance [50]; combined with the Mandelstam–Tamm form, the fastest attainable evolution is set by whichever of mean and spread is smaller [23]. Lloyd’s rephrasing — that a system of average energy E performs at most $2E/\pi\hbar$ elementary operations per second [46] — is the same statement counted as a rate, and it is the form the foam lineage later leans on. The geometric picture underneath all of these, that distinguishability is a distance on the manifold of quantum states measured by the Fisher information, is Braunstein and Caves’ [13]; it is what lets the metrology bounds be stated as scalings rather than case-by-case results. Those scalings are the second line of Proposition 1: repeating a measurement N times, or running N atoms in parallel, reduces statistical error as $N^{-1/2}$, while entanglement can push the reduction to N^{-1} [32]. The standard quantum limit and the Heisenberg limit bracket what a clock’s instability can do with a fixed number of atoms and a fixed averaging time; we will use both as brackets, not commit to either, because which one a given clock reaches is an engineering fact and the floor of this paper is not.

What the speed limits leave open is the question the rest of the paper turns on: *whose* energy is being counted. Equation (1) bounds the spread of the oscillating degree of freedom. A real clock is more than its oscillator — it has a mass, a mount, a laser, a vacuum can — and the speed limits say nothing about which of that mass-energy is “the clock’s own” in the sense that will gravitate and fluctuate together with the tick. The quantum-clock resource theory is where that question acquires an answer, and this is what it adds beyond the bare limits. By treating the clock as an autonomous machine that runs on its own free energy rather than an externally driven oscillator, the resource theory forces one to say what the clock *is* as a closed system. Ranković, Liang, and Renner’s alternate-ticks game makes “how good is this clock” an operational quantity — how long two copies

can be kept ticking in step [76] — so that the clock’s quality is defined without reference to an external standard, which is exactly the setting in which an irreducible self-effect can even be posed. Woods, Silva, and Oppenheim show that a finite quasi-ideal clock can drive a target system with a back-reaction that falls exponentially in the clock’s Hilbert-space dimension while its energy grows only linearly [95], and Woods’ later axiomatic construction unifies the autonomous and high-accuracy models into one [94]; together they establish that one can, in principle, make the clock’s coupling to everything except its own reading arbitrarily clean — which is what licenses us, in Section 3, to isolate the energy that fluctuates *conditionally on the tick* as a well-defined subsystem rather than a hand-wave.

The resource theory also supplies a second, thermodynamic floor that runs parallel to ours and must not be confused with it. Erker and collaborators showed that any autonomous quantum clock produces entropy, and that its accuracy is bounded by the heat it dissipates [29]; Meier and collaborators proved a universal trade-off between resolution and accuracy for clocks whose elementary events are memoryless, fixed by the number of stochastic events per reference tick [55]; Schwarzhans and collaborators showed a clockwork can be made arbitrarily accurate only by growing its complexity, with the second law as the price [83]. Milburn frames the whole subject by noting that every clock, classical or quantum, is an open irreversible system whose central limiting mechanism is phase diffusion [57], and the single-qubit case makes the trade-off quantitative: the energy cost of a clock is at least $k_B T$ times the entropy reduction it achieves per reading, so quantum mechanics forbids exact timekeeping even in principle for a one-qubit clock [33]. That this thermodynamic trade-off is real and near has been measured: a nanomechanical clock was shown to obey a linear accuracy-versus-entropy relation and to run within an order of magnitude of its theoretical bound [67], and imperfect timekeeping has been shown to cap the achievable complexity of quantum computation [56, 96], with related dimension-versus-sharpness bounds on any generated time signal [86, 97]. We keep this thermodynamic price in the ledger (Fig. 1) because it is a genuine quantum charge, but we mark the distinction the accounting depends on: the entropy floor bounds a clock’s accuracy through its *environment*, the heat it must shed, whereas the floor this paper derives bounds it through the clock’s *own gravity*, and the two are independent columns. The resource theory’s real gift to Section 3 is not this second bound but the first thing it forced us to fix — a defensible answer to whose energy the clock’s rate depends on — because that is the energy general relativity is about to charge rent on.




Figure 1: The two prices a clock pays (Propositions 1 and 2). Quantum mechanics charges rate resolution in energy spread σ_E (left); general relativity charges the energy a clock contains a self-redshift (right). The paper’s contribution is the split on the right: the *mean* self-redshift is a calibrable systematic and sets no limit, while only the *fluctuation* of the gravitating energy is irreducible — and that fluctuation is bounded below by the very quantum price on the left. The two meet in the accuracy floor $\mathcal{A}_{\min}$, quadratic in the Planck time rather than the $\tau^{1/3}$ of the spacetime-foam lineage.

3 The clock’s own gravity charges rent, but only the fluctuating part of the bill cannot be paid off

The quantum price of Section 2 is charged before gravity enters. Gravity’s charge is separate and has a structure the heuristic lineage never resolved, because it requires distinguishing two things that the phrase “the clock’s energy” runs together. A clock contains mass-energy; that mass-energy sources a potential; the clock’s rate is redshifted in the potential it makes. So far this is arithmetic every party agrees on. The question the arithmetic hides is whether that self-redshift limits how good the clock can be, and the answer splits cleanly the moment one asks which part of the energy is a fixed number and which part is a quantum variable correlated with the reading. That split is this paper’s contribution.

Proposition 2 (The relativistic rent and its irreducible part; ours, exact in the model). *Model a clock as an oscillator of size R carrying mean rest energy $\bar{E}$ with quantum fluctuation σ_E , where σ_E is bounded below by Proposition 1 for the resolution the clock claims. The clock’s rate is redshifted in its own gravitational potential,*

$$f \longrightarrow f \left(1 - \frac{G \bar{E}}{R c^4} \chi \right), \quad (2)$$

with χ an order-unity geometry factor set by the mass distribution. The mean shift is a fixed systematic: known

once the geometry is known, and therefore calibrable in principle and no bound on accuracy. The fluctuation is not calibrable, because it is correlated with the tick record itself; it dephases the clock against any external standard by

$$\frac{\delta f}{f} \sim \frac{G \chi \sigma_E}{R c^4}. \quad (3)$$

Since Eq. (3) increases in σ_E while Proposition 1 bounds σ_E below, the least rent is paid at the least energy spread the resolution allows, giving an irreducible contribution to fractional accuracy

$$\mathcal{A}(R, \tau) \geq \frac{G \hbar \chi}{R c^4} \cdot \frac{\delta \phi}{2 \delta t}. \quad (4)$$

Which subsystem’s σ_E enters Eq. (3) — the oscillator alone, or the clock’s total gravitating mass-energy — is bounded both ways in Appendix A, and the two bounds bracket the floor.

Equation (2) is the weak-field redshift of a clock in the potential of its own energy $\bar{E}$ at scale R ; its derivation from the metric, and the fixing of χ for a uniform sphere and for a shell, are in Appendix A. We grade the equation exact in the model and specify the model rather than defend it in the abstract: a localized non-relativistic clock, treated to first post-Newtonian order, with its self-potential evaluated at its own centre. Nothing in this section needs χ to more than order unity, and nothing needs the clock to be spherical; what the argument needs is the one distinction the display makes visible, between $\bar{E}$ and σ_E .

The mean shift is genuinely large and genuinely calibrable, and both halves of that sentence matter. It is large: an optical clock already resolves the difference in redshift across a millimetre of the Earth’s field, the JILA measurement reading a 7.6×10^{-21} gradient across a millimetre-scale strontium sample [11], and clocks compared over a continent are used the other way, as gravimeters, resolving centimetre height differences from the redshift alone at the 10^{-18} level [53, 54]. The self-redshift of the clock’s own mass is the same effect turned inward, and it is far larger than any accuracy the clock claims. And yet it does not bound accuracy, because it is a fixed number. Chronometric geodesy is the existence proof: the entire enterprise of comparing clocks across national metrology institutes rests on computing and subtracting the relativistic redshift correction from the known gravitational potential [24], and the same subtraction applies to the potential the clock makes for itself. A large systematic that you know is not a limit on accuracy; it is a term you carry. This is the point the heuristic lineage passes over. Salecker and Wigner required a minimum clock mass so the clock’s quantum spread would not blur the interval [79], and read that minimum mass as a cost; Ng and van Dam propagated the mass through a light-crossing and a Schwarzschild-radius constraint to a length uncertainty $\delta l \propto l^{1/3}$ [60, 61]; Lloyd

counted the total energy $\bar{E}$ as the operative resource [46]. In each the quantity doing the bounding is the mean — m , $\bar{E}$, the total — and the mean is the calibrable part. We claim, and Section 6 argues in detail, that this conflation is why those bounds overstate what the self-effect forbids.

What cannot be calibrated away is the part of the redshift that is a quantum variable correlated with the reading, and identifying it precisely is the technical move. The clock’s rate depends on its own energy through Eq. (2); if that energy is in a superposition, or carries the irreducible spread σ_E that Proposition 1 demands, then the clock’s rate is itself uncertain, and the uncertainty is not independent of the tick — it is the tick’s own energy fluctuation seen gravitationally. Castro-Ruiz, Giacomini, and Brukner made the mechanism explicit: a clock’s mass-energy superposition sources a non-fixed metric, so that two clocks are entangled through their mutual time dilation and “the time read by a single clock is not well defined” at the quantum level, classical proper time returning only in the decohered limit [18]. That is the cross term of Eq. (3) named in its own right. Pikovski and collaborators showed the same coupling as a decoherence: gravitational time dilation dephases a composite body through the fluctuation of its internal energy, with a decoherence time $\tau_{\text{dec}} \approx \sqrt{2/N} \hbar c^2 / (k_B T g \Delta x)$ for N constituents in a superposition of extent Δx [73]. The basis matters and it is the right one: Anastopoulos and Hu derive, from quantum field theory on a weakly curved background with no added assumptions, that gravitational decoherence acts in the *energy* basis rather than the position basis [5] — which is exactly the statement that it is the energy fluctuation σ_E , not a position spread, that the self-redshift reads out. We grade the identification of Eq. (3)’s σ_E with the conditionally-fluctuating internal energy as structured: it is a leading-order reading of these mechanisms, and the coefficient hidden in “ $\sim$ ” is where the model does its work, which is why Appendix A carries the two bracketing derivations rather than a single number.

The distinction survives the obvious objection, which is that the mean is also uncertain — one never knows $\bar{E}$ exactly, so is the “calibrable” mean not itself a source of error? The answer is the operational one this series has used throughout: the mean is a c-number one measures against an external standard and subtracts, its residual uncertainty falling with measurement effort like any systematic, while the fluctuation σ_E is the part that remains correlated with *this* tick and cannot be subtracted by any amount of prior characterization, because the characterization would have to know the outcome of the tick it is trying to time. That this is a real distinction and not a verbal one is shown by how sharply the community disputes what a redshift measurement even fixes. When an atom interferometer was presented as measuring the gravitational redshift at the Compton frequency to 7×10^{-9} [58], Wolf and collaborators argued

the Compton-phase difference between the arms is zero in general relativity and a large class of theories, so the experiment tests the universality of free fall rather than the redshift [93]; Hohensee and collaborators replied that atom interferometers function as redshift tests in any formalism [36]. We take no side in that exchange here — both cases are strong — but we take its lesson, that a relativistic reading fixes a definite quantity only once one says which one, and apply it to the clock’s own energy: the calibrable quantity is the mean shift of Eq. (2), the irreducible quantity is the fluctuation of Eq. (3), and no experiment confuses them once they are written down.

Two further threads confirm the split and one bounds how close a laboratory already is. Zych and collaborators’ interferometric-visibility witness makes the internal-energy fluctuation observable as a loss of coherence when a clock is put in superposition of proper times [100], and Khandelwal, Lock, and Woods show that a clock with a classical motional state recovers ordinary general-relativistic time dilation exactly, while a delocalized clock suffers a genuine, motion-state-dependent discrepancy between proper time and the time it reads out [39] — the discrepancy being, again, the irreducible part. Smith and Ahmadi derive the interaction that produces it, a clock-system coupling of order G/c^4 inversely proportional to their separation [85], which is the same G/c^4 that scales Eq. (3). How close is the laboratory? Paczos and collaborators estimate the quantum correction to time dilation for a clock in superposition of heights at $\sim 10^{-20}$ for a millimetre separation, flagging that scale as within reach of the best atomic clocks [64], and Roura’s clock-interferometry proposal identifies ten-metre atomic fountains as the apparatus that would isolate the genuinely general-relativistic signal [78]. These fix the observable near the top of the map; the floor lies far below, and Section 4 works out how far.

The split has a genuine boundary, and it is the sharpest objection to the whole construction: is the mean even well-defined? Castro-Ruiz and collaborators’ own result is that a mass-energy superposition sources a metric that is itself in superposition [18], so “the mean shift” of Eq. (2) is a semiclassical notion — the redshift in the metric sourced by the expectation value $\bar{E}$ — not an exact one. We therefore grade the calibrability of the mean structured rather than exact: it holds to the order at which the clock’s own field is a fixed background sourced by $\bar{E}$, which is the weak-field, non-relativistic regime every real clock occupies by many orders of magnitude. The claim would fail where the clock’s energy superposition is wide enough that no single background metric approximates it, and no clock approaches that regime. What would kill the split operationally is more concrete: a measured residual self-redshift, after full geodetic calibration against an external standard, exceeding the fluctuation bound of Eq. (3) — a mean that provably cannot be subtracted. Until such a residual is shown, the mean is a term one carries and the fluctuation

is the floor.

A word on what this section does *not* claim: the self-energy criterion of Diósi and Penrose, in which a mass superposition collapses on a timescale $\hbar/E_G$ set by its gravitational self-energy [26, 27, 71], is a different proposal — a modification of quantum dynamics — and our rent is not a collapse. We note the resemblance of the self-energy bookkeeping and grade the analogy speculative: Eq. (3) is a dephasing within standard semiclassical gravity, requiring no new physics, and it is the conservative reading of the same coupling those authors read maximally. The two columns of the ledger — quantum price and relativistic rent — are now both on the table (Fig. 1); combining them is the next section’s arithmetic.

4 The rent falls with size and the price falls with energy, and a coherent tick picks the clock that pays least

The two columns of the ledger pull in opposite directions on the same knob. The relativistic rent of Eq. (4) falls as $1/R$: a larger clock dilutes its self-potential and pays less. The quantum price falls the other way, with more energy and more atoms sharpening the tick and lowering the instability as $1/(\sqrt{N}\tau)$ or better [32]. Left alone, both would send the accuracy to zero — build the clock bigger and drive it harder. What forbids that, and turns the trade into a finite floor with a definite best size, is a constraint the heuristic lineage supplied as its punchline and we use as a mere boundary condition: a tick is one synchronized event only if the clock is no larger than the distance light crosses during the tick.

Proposition 3 (Scaling, the light-crossing constraint, and the optimal clock; ours). *The rent of Eq. (4) decreases monotonically in R , so accuracy improves with size. A coherent tick requires that the whole clock share one reading, which bounds the size by the light-crossing distance during the tick,*

$$R \leq c\delta t. \quad (5)$$

Minimizing Eq. (4) over R subject to Eq. (5) places the optimum at the boundary, $R^* = c\delta t$, and yields the absolute floor

$$\mathcal{A}_{\min}(\tau) \geq \frac{G\hbar\chi\delta\phi}{2c^5\delta t^2} = \frac{\chi\delta\phi}{2} \left(\frac{t_P}{\delta t}\right)^2, \quad (6)$$

where the second form uses $t_P^2 = \hbar G/c^5$ and δt is set by the averaging time τ and the demanded phase resolution $\delta\phi$. The floor is quadratic in the Planck time, not the one-third power of the foam lineage.

The light-crossing constraint (5) is the Salecker–Wigner ingredient [79], and stating it as a constraint rather than a result is the whole difference in how it is used. Salecker and Wigner, and Ng and van Dam after them, ran the light-crossing bound forward together with a minimum mass to produce an intrinsic length uncertainty; we run it as the one thing that stops the rent from being lowered without limit. A clock larger than $c\delta t$ cannot register a single tick, because its far side has not learned that its near side has ticked; so $R = c\delta t$ is the largest coherent clock, and since the rent falls with size, the largest coherent clock is the best one. There is no interior optimum to find — accuracy improves right up to the coherence boundary and the boundary is the answer. This is worth saying plainly because the lineage’s phrasing suggests a mysterious optimal scale; there is nothing mysterious in it. The best clock is as big as it can be while still ticking as one thing.

Substituting $R^* = c\delta t$ into the rent gives Eq. (6), and the second form is the one to read. The floor is $(\chi\delta\phi/2)$ times $(t_P/\delta t)^2$: the square of the ratio of the Planck time to the clock’s resolution time. Two features are worth grading. First, the scaling is exact in the model — it follows from Eqs. (4) and (5) with no further assumption — and it is a genuinely different scaling from the lineage’s. The Ng–van Dam clock bound reads $\delta t \gtrsim t_P^{2/3}\tau^{1/3}$ [61], a one-third power; ours is a square. The difference is not a coefficient dispute but a structural one, and its origin is exactly the accounting of Section 3: the lineage’s $\tau^{1/3}$ comes from charging the mean energy, ours from charging only the fluctuation, and the two energies scale differently with the clock’s parameters. Section 6 returns to this as the core of the graded verdict. Second, the coefficient is honestly order-unity-uncertain: χ is a geometry factor, $\delta\phi$ a design choice, and the bracketing of Appendix A moves the whole floor by the ratio of oscillator energy to total clock energy. None of that touches the exponent, and the exponent is the point.

The averaging time enters through δt , and tracing it gives the floor’s τ -scaling and, with it, two regimes worth separating. The resolution time δt is the interval over which a single distinguishable tick forms; it cannot exceed the total averaging time, since a tick that has not completed within τ has not been read, so $\delta t \leq \tau$. The floor Eq. (6) decreases in δt , so the best floor a clock of unconstrained size can reach uses $\delta t = \tau$, giving

$$\mathcal{A}_{\min}(\tau) \geq \frac{\chi\delta\phi}{2} \left(\frac{t_P}{\tau}\right)^2, \quad R^* = c\tau, \quad (7)$$

an absolute floor that improves quadratically with averaging time and an optimal size that grows with it — a clock as large as light crosses in the whole averaging interval. That is the ideal regime, and it is not the laboratory’s. A real clock’s size is capped far below $c\tau$ by the bench it sits on, and then the binding constraint is not the averaging time but the fixed size: with R pinned at a metre, $\delta t = R/c$ is a few nanoseconds regardless of how

long one averages, the light-crossing constraint Eq. (5) saturates at the hardware scale, and the floor sits at the size-limited value of Eq. (4) rather than the time-limited value of Eq. (7). The optimal-size corollary is thus read two ways on the same plot: along the locus $R^* = c\delta t$ a clock does the best its coherence allows, and where that locus is truncated by the physical size of the apparatus is where a given instrument actually lives. The gap between the two — between a laboratory clock’s floor and the ideal τ -limited floor a clock of size $c\tau$ would reach — is itself many orders, and it is why the wall is even farther below the bench than below the ideal.

The exponent is what makes the floor a wall no one is near rather than a target someone is approaching. Evaluated at laboratory scale — a clock of order a metre, so δt of order the light-crossing time R/c — the ratio $t_P/\delta t$ is so small that its square lands the floor scores of orders of magnitude below the demonstrated frontier. We quote the landing from the seeded computation of Appendix C, which fixes χ , $\delta\phi$, and the bracket rather than leaving them as symbols, and we quote it with order-honesty: the figure that matters is the power of ten, and the coefficient is not claimed to more than that. At $R = 1$ m, with $\delta t = R/c = 3.34 \times 10^{-9}$ s and $t_P = 5.39 \times 10^{-44}$ s, the floor lands at $\mathcal{A}_{\min} \approx 1.3 \times 10^{-70}$ (Appendix C) — about 52 orders of magnitude below the demonstrated 9.4×10^{-19} frontier. The floor family over R , with the optimal-size locus $R^* = c\delta t$ marked, is drawn in Fig. 2; the locus is the line along which a clock of each size does the best its coherence allows, and the absolute floor is where that line meets the largest size a given δt permits. The floor falls as R^{-2} along that locus (Fig. 2), so the optimal clock is the largest coherent one the apparatus allows, and the metre-scale bench sets the laboratory floor near 10^{-70} .

Three honest boundaries close the section. The first is what the floor is a floor *on*: it bounds the accuracy achievable against an external inertial standard by a clock whose own gravity is the only uncalibrated effect, and it says nothing about a clock read against a copy of itself in the same potential, where the self-redshift is common-mode and cancels — which is why the alternate-ticks operational quality of two co-located clocks [76] is not what Eq. (6) bounds. The floor is a limit on *absolute* accuracy, on knowing one’s rate against the flat-space ideal, not on relative stability. The second boundary is the model’s: Eq. (6) is first-post-Newtonian and semiclassical, and it assumes the fluctuating energy gravitates linearly; a full quantum-gravitational treatment could change the coefficient or, if spacetime genuinely fluctuates, add a term, which is the possibility Section 6 weighs. The third is measurement: the floor would be falsified as a floor if any clock were shown to beat Eq. (6) against an inertial standard, and it would be promoted from a consistency bound to a discovery if a clock resolved the fluctuation rent of Eq. (3) itself — a departure from the calibrable mean at the $(t_P/\delta t)^2$ level. Neither




Figure 2: The accuracy floor over clock size (Proposition 3). Along the optimal-size locus $R^* = c\delta t$ the floor falls as R^{-2} , so a larger coherent clock does better, bounded only by the largest size the apparatus permits. At laboratory scale ($R \sim 1$ m) the floor is $\mathcal{A}_{\min} \approx 10^{-70}$ — about 52 orders of magnitude below the demonstrated 10^{-19} frontier (red). The floor is a wall no one is near, not a target anyone is approaching; the exponent, not the coefficient, is the claim (Appendix C fixes χ , $\delta\phi$, and the bracket).

is close, and the size of the gap between the frontier and the floor, drawn in the next section on the full landscape, is the quantitative form of that statement.

5 Ten decades of won frontier and the headroom beneath it fit on one plane

A floor is only as interesting as its distance from where instruments actually sit, and that distance is best read off a single plane. This section places the demonstrated clocks, the near-term projections, and the floor family of Section 4 on the plane of clock size R against fractional accuracy $\mathcal{A}$. It is the series-closing figure: ten papers of operational quantum-relativity ending on the full landscape of how good a clock is, how good one might soon be, and how far either is from the wall.

Proposition 4 (The landscape; ours as compilation). *On the plane of clock size R against fractional accuracy $\mathcal{A}$, the demonstrated instruments, the near-term projections, and the floor family of Proposition 3 occupy separated regions. The flagship optical clocks cluster near 10^{-18} – 10^{-19} ; the projected nuclear and entanglement-enhanced clocks extend the frontier by at most a few orders; the floor lies far below all of them at every laboratory size. Table 1 lists each plotted instrument with*

Table 1: The instruments of Fig. 3 (Proposition 4): representative clocks, their approximate size, and demonstrated or projected fractional accuracy. Even the best sit some fifty orders above the floor of Fig. 2.

clock	R [m]	accuracy	source
Essen–Parry Cs (1955)	~ 1	10^{-10}	[30]
Cs fountain	~ 1	10^{-16}	[48]
Sr optical lattice	~ 0.3	8.1×10^{-19}	[2]
Al^+ single ion	~ 0.3	9.4×10^{-19}	[14]
Th-229 nuclear (proj.)	~ 0.1	$\sim 10^{-19}$	[17]
GPS block clock	~ 0.3	10^{-13}	[48]

its size, accuracy, averaging time, and source key; Fig. 3 draws them against the floor.

The frontier is a genuine achievement and worth recounting in numbers, because its slope over seventy years is what sets the scale for reading the gap. It begins with Essen and Parry’s caesium resonator, the first atomic frequency standard [30], built on Ramsey’s separated-oscillatory-field interrogation [75] that still underlies every microwave and optical clock. Caesium fountains realizing the SI second reach fractional accuracy approaching 10^{-16} [10]; the move to optical transitions, some five orders of magnitude higher in frequency than the caesium microwave [48], opened the two decades of gain that followed. The single-ion and lattice frontiers now sit together near the eighteenth and nineteenth digits: $^{27}\text{Al}^+$ quantum-logic clocks at 9.4×10^{-19} [14] and, in the direct successor instrument, 5.5×10^{-19} with a stability of $3.5 \times 10^{-16}/\sqrt{\tau/\text{s}}$ [51]; ^{87}Sr optical lattices at 8.1×10^{-19} [2]; ytterbium lattices at 1.4×10^{-18} with measurement instability 3.2×10^{-19} [53]; a $^{171}\text{Yb}^+$ single-ion octupole clock at 3×10^{-18} [38], cross-validated against a second such clock at the 10^{-18} level over six months [80]; and cryogenic ^{87}Sr lattices at the low 10^{-18} level, two of which agreed to 2.0×10^{-18} [90]. The three-clock network comparison of Al^+ , Sr, and Yb reached ratio uncertainties of 8×10^{-18} [6], and the JILA millimetre-gradient measurement resolved a fractional frequency of 7.6×10^{-21} across a millimetre [11], with a multiplexed lattice reaching a statistical uncertainty of 8.9×10^{-20} after 3.3 hours [99]. These are the points near the top of Fig. 3.

The clocks also occupy a wide range of sizes, and the map’s horizontal axis is not decoration. At the compact end, a chip-scale atomic-beam clock reaches $1.2 \times 10^{-9}/\sqrt{\tau}$ [52] and a chip-scale atomic clock holds timing through twelve hours of GPS outage to within four microseconds [66], the GPS/telecom-grade anchor for scale; in the middle sit transportable optical lattices, a ^{87}Sr clock at 7.4×10^{-17} in a trailer [40], a transportable Sr clock with a 4.0×10^{-19} blackbody-radiation-shift uncertainty and 2.1×10^{-18} total [62], a 35-litre optical clock run at sea below 10^{-18} in the laboratory [77], and a Yb transportable lattice shipped 3000 km and opera-

tional in two days [12]; and at the large end sit the fibre networks that make separated clocks one instrument, 1415 km linking Paris and Braunschweig at 5×10^{-17} [45] and a 2220 km link at 7×10^{-17} instability [81]. Plotting size against accuracy is what lets the floor family of Section 4, itself a function of R , be drawn through the same instruments rather than beside them.

The near-term projections extend the frontier but not the picture. The nuclear clock is the most concrete, and its energy has been pinned down over two decades in a way worth recounting because the convergence is what makes it a plottable point rather than a hope: proposed at a mistaken 3.5 eV [70], first directly detected in a wide 6.3–18.3 eV bracket [91], fixed by internal-conversion spectroscopy at 8.28(17) eV [84], seen in radiative decay at 8.338(24) eV with a 670(102) s host lifetime [42], and finally resolved by direct laser excitation at 8.355733(2) eV [28], corroborated at 8.4 eV in an independent host [88] and linked directly in frequency to a strontium clock [98]. Its projected performance is still far from the atomic frontier — a solid-state design reproduces its frequency between crystals only to 1.4×10^{-13} today, though the same work projects suppression of the leading systematic below 10^{-18} [63], and a feedback-stabilized design projects approaching 10^{-15} after a day’s averaging [89] — but the single-ion nuclear-clock proposal targets 10^{-19} [17], so it belongs on the map as a projected point near the current frontier, not beyond it. Its real promise is elsewhere, an enhanced sensitivity to variation of the fine-structure constant with an enhancement factor $K = 5900(2300)$ [8, 9, 69, 87]; that is a reason to build it, but not a lower accuracy floor. The entanglement route is similar. A network of clocks sharing entanglement can, in principle, operate near the Heisenberg limit [41], and entanglement on an optical clock transition has delivered a 4.4 dB metrological gain, a factor 2.8 faster averaging than the standard quantum limit [68]; but the gain is a stability speed-up, not a new accuracy regime, and it is bounded — spin squeezing helps a strontium lattice only below about a thousand atoms at present laser noise [82]. The redefinition of the second is being organized around reaching 10^{-18} across the supporting infrastructure [25], which is the official statement of where the frontier is headed: another order or two, not another forty.

One band of the map deserves marking because it is where the thesis of this paper is already visible in the instruments. The best clocks have entered the regime where the gravitational redshift is not a correction but a resolved signal: a ytterbium lattice resolves a one-centimetre height difference through the redshift at the 10^{-18} level [53], the JILA measurement reads the redshift gradient across a single millimetre-tall sample [11], and Roura’s proposal identifies clocks already sensitive to a one-centimetre height difference, with ten-metre fountains as the next step [54, 78]. Every one of these measures the *mean* redshift — the calibrable systematic




Figure 3: The clock landscape (Proposition 4), closing the series. Every flagship timepiece from Essen and Parry’s first caesium standard (10^{-10} , 1955) to today’s optical and projected nuclear clocks ($\sim 10^{-19}$) sits in a narrow band near the top; the accuracy floor $\mathcal{A}_{\min}(R)$ that quantum mechanics and general relativity jointly permit runs scores of orders below it. The double arrow marks the gap at laboratory scale: about 52 orders of magnitude of accuracy the two theories leave on the table and no instrument has used. Values and floor from the seeded computation of Appendix C; instrument accuracies from the references of Table 1.

of Section 3 — and does so exquisitely. None measures the irreducible fluctuation, which is the floor. That the frontier now sits squarely on the calibrable term, forty orders above the irreducible one, is the single most important thing the map shows, and it is the picture the closing section reads against the foam claims.

That is the shape of the landscape, and the shape is the argument. The frontier has moved roughly ten orders of magnitude from Essen’s caesium to today’s optical clocks, and the most optimistic near-term projections add a few more. The floor of Section 4 lies scores of orders below even those, and the vertical gap between the projected frontier and the floor at laboratory size — the empty middle of Fig. 3 — is the quantitative content of “finite but far.” That empty middle is about 52 orders of magnitude tall at metre scale (Appendix C): the accuracy the two theories jointly permit and no clock, from Essen and Parry’s 10^{-10} to today’s 10^{-19} , has come within fifty orders of spending. Every instrument plotted carries its verified number and source in Table 1; the compilation is ours only in its assembly and its placement against the floor, and we grade it exactly that way — no number here is derived, each is measured and cited, and the one thing the plate asserts on its own authority is that the floor runs beneath all of them with room to spare. What that room means for the Planck-foam reading of the same picture is the closing section’s subject.

6 The foam bounds charge the bill we have just shown is calibrable, and the floor is a consistency bound, not a discovery

The floor of Section 4 sits scores of orders below the frontier, and the temptation — old, and repeatedly voiced in the lineage of Section 1 — is to read that vast headroom as a promise: a clock good enough would be reading the quantum structure of spacetime, the Planck-scale foam, off its own instability. This section grades that reading. It reconstructs the foam clock bounds honestly, names the counting assumption each rests on, applies the accounting objection of Section 3, sets the published data against them, and then says plainly what our own floor is and is not — which is where the series ends.

Proposition 5 (The foam verdict; ours as critique, graded per claim). *The holographic clock bounds — Ng–van Dam’s $\delta t \gtrsim t_P^{2/3} \tau^{1/3}$ and its relatives — follow from counting the clock’s total mass-energy as the resource that gravitates, exactly the mean quantity Proposition 2 shows is calibrable. Reconstructed in Appendix B, they conflate the calibrable mean with the irreducible fluctuation; graded as a structured objection, this places the $\tau^{1/3}$ scaling among conjectures rather than consequences. The astrophysical wavefront-coherence data exclude the corresponding length scaling $\alpha = 2/3$ directly, and the floor of Proposition 3 is a consistency bound within semiclassical gravity, not a measurement of quantum spacetime.*

The bounds deserve to be stated at their strongest be-

fore they are graded. Ng and van Dam’s runs as follows: to measure a distance l one needs a clock at each end; the clock must be massive enough that its own quantum spread stays small over the light-crossing time (the Salecker–Wigner requirement [79]), but not so massive that it collapses inside its own Schwarzschild radius; optimizing the mass between these gives an intrinsic length uncertainty $\delta l \sim (l \ell_P^2)^{1/3}$, equivalently $\delta l/l \sim (\ell_P/l)^{2/3}$ [60, 61], and read as a clock the same optimization gives $\delta t \gtrsim t_P^{2/3} \tau^{1/3}$. The holographic reading makes it look inevitable: the number of degrees of freedom in a region is bounded by its area in Planck units, and Ng turned that into a limit tying a computer’s speed and memory, $\tilde{\nu}^2 I^{-1} \lesssim t_P^{-2}$, validated against black-hole evaporation [59], with Lloyd’s ultimate-computation limits — $2\bar{E}/\pi\hbar$ operations per second, 10^{120} operations over the history of the universe — as the same accounting at cosmic scale [46, 47]. Reconstructed in full in Appendix B, this is a real and coherent line of reasoning, and we report it as the motivation it was, not a strawman.

It has, however, two published objections and now a third. The first is measurement-scheme dependence, shown by Baez and Olson: replace the free-falling point mass with a massive elastic rod and the $L^{1/3}$ enhancement disappears, the uncertainty falling to the bare Planck length ℓ_P set by the rod’s zero-point motion, because a rigid extended body evades the optimization that produced the enhancement [7]. Calmet, Graesser, and Hsu reach a Planck-length floor from ordinary quantum mechanics and *classical* general relativity with no $L^{1/3}$ factor at all [16], and Padmanabhan had earlier argued the bare Planck length and Planck time as the operational floors on position and synchronization [65]; Garay’s review and Hossenfelder’s catalog both present the minimal-length scale as model-dependent, with the escape routes a recurring feature rather than an accident [31, 37]. So the $\tau^{1/3}$ scaling is, on the theory side, one derivation among several that reach different exponents, and its enhancement over the bare Planck scale is the disputed part. Our own objection is the third, and we grade it structured: the Ng–van Dam optimization charges the clock’s *mean* mass, the quantity Section 3 shows is a calibrable systematic. A bound built on a quantity one can measure and subtract does not bound accuracy; it bounds a clock that has not been calibrated. When only the irreducible fluctuation is charged, the exponent changes from one-third to the square of Eq. (6), which is why our floor lies far below the foam bound rather than at it — the two are computing charges on different energies.

The data settle the length form directly, and this is the strongest part of the verdict. Perlman and collaborators parameterize foam models by an accumulation power α via $\delta\ell \simeq \ell^{1-\alpha} \ell_P^\alpha$, the random-walk model at $\alpha = 1/2$ and the Ng–van Dam holographic model at $\alpha = 2/3$, and test them against the preservation of wavefront coherence in images of distant sources: Chandra X-ray data require

$\alpha \gtrsim 0.58$, Fermi GeV data $\alpha \gtrsim 0.67$, and ground-based TeV data $\alpha \gtrsim 0.72$, all three exceeding $2/3$, so the data appear to exclude $\alpha = 2/3$ [72]. We grade this claim strong and largely dispositive for the length scaling: the holographic foam model in the form the lineage made famous is excluded by direct astrophysical measurement, independent of any accounting argument. The interferometric side agrees within its own mechanism: Hogan predicted a holographic “shear” noise with amplitude $\sim \sqrt{ct_P L}$ on the light-crossing time of a region of size L , with no free parameters [34, 35], and the Holometer — two co-located 40 m interferometers cross-correlated past their 7.6 MHz light-crossing frequency — returned a null, ruling out that shear-noise class by a large factor [19, 43]. A caution belongs here, and we grade the older phase-coherence bounds contested: Lieu and Hillman argued from Hubble Space Telescope diffraction rings of a 1.2 Gpc galaxy that first-order foam with amplitude $a_0 > 0.003$ is excluded, and hence that the speed of light is exact to a few parts in 10^{32} [44], but whether unresolved-source imaging isolates Planck-scale decoherence from ordinary point-spread systematics has been questioned, and a companion study reaching a compatible conclusion by the same method on different targets [74] may be independent confirmation rather than the rebuttal it is sometimes filed as; we report the phase-coherence bound as claimed and contested, not as settled. And one adjacent constraint must be kept separate: the 31 GeV photon from GRB 090510 arriving within a second, which bounds any linear energy-dependent photon dispersion to above the Planck mass [1], tests a *dispersion* signature of foam, not the position-noise or phase-noise signatures the other bounds address, and conflating the three is a common error we decline.

What, then, is our floor? It is a consistency bound within semiclassical gravity, and calling it anything more would repeat the mistake we have just graded. Equation (6) follows from the weak-field redshift of a clock in its own potential and the Mandelstam–Tamm price of its own tick; it invokes no quantized geometry, and it would hold in a universe with no foam at all. It says that a clock read against an inertial standard cannot do better than $(t_P/\delta t)^2$ because its own fluctuating mass-energy dephases it, and that is a statement about the clock, not about spacetime. The Planck time appears in it for dimensional reasons — $t_P^2 = \hbar G/c^5$ is just the combination of the two charges’ constants — not because the clock is probing Planck-scale physics.

That distinction is the right note for the series to end on, because it says exactly what would count as more. A clock would be reading spacetime, rather than calibrating itself, if it resolved a departure from the calibrable mean of Eq. (2) — either the irreducible fluctuation rent of Eq. (3) itself, the genuinely quantum correction to time dilation that Castro-Ruiz and collaborators and Pikovski and collaborators identified [18, 73], or a foam term riding above our floor. The first of these is not fan-

tastically far: Paczos and collaborators put the quantum time-dilation correction for a millimetre height superposition at $\sim 10^{-20}$ [64], and ten-metre atomic fountains are the identified apparatus [78] — so the discovery that matters is not a clock forty orders better, but a clock that puts its *own energy* in superposition and reads the difference. This is where the arc of these papers has been heading. The operational treatment of self-redshift and of proper time as a measured quantity was built in the series’ opening papers [20, 21], and the discipline of asking what quantity a relativistic reading actually fixes — the discipline that let us separate the mean from the fluctuation here — was developed across the middle of the series [22]. Ten papers of operational quantum-relativity close on the same lesson their subject keeps teaching: a clock this good is not a window onto quantum spacetime but a mirror held up to its own energy, and telling the two apart is the whole of the measurement.

A Derivation of the rent, the floor, and the two brackets

This appendix derives Eqs. (2)–(6) of the main text, fixes the geometry factor χ for two mass distributions, and states the two bracketing bounds on which subsystem’s energy fluctuation enters the rent. The register is that of a semiclassical, first-post-Newtonian calculation; every place a stronger assumption is smuggled in is marked.

The self-redshift. A clock at rest at position $\mathbf{r}$ in a static weak field with Newtonian potential $\Phi(\mathbf{r})$ runs slow relative to a distant inertial standard by the standard redshift factor

$$\frac{d\tau}{dt} = \sqrt{1 + \frac{2\Phi}{c^2}} \approx 1 + \frac{\Phi}{c^2}, \quad \Phi < 0, \quad (8)$$

so a rate f is observed shifted by $\delta f/f = \Phi/c^2$. For a clock that is the source of its own field, Φ is the self-potential of its mass-energy $\bar{E} = Mc^2$ evaluated at the clock. Writing the self-potential at scale R as $\Phi_{\text{self}} = -\chi GM/R = -\chi G\bar{E}/(Rc^2)$ with χ a dimensionless geometry factor gives

$$\frac{\delta f}{f} = \frac{\Phi_{\text{self}}}{c^2} = -\frac{G\bar{E}}{Rc^4} \chi, \quad (9)$$

which is Eq. (2). The factor χ is the only geometry input, and it is order unity. For a uniform sphere of radius R the potential at the centre is $-\frac{3}{2}GM/R$, so $\chi = 3/2$; the potential averaged over the volume is $-\frac{6}{5}GM/R$, so a volume-averaged reading gives $\chi = 6/5$; a thin shell of radius R has interior potential $-GM/R$, so $\chi = 1$. Nothing in the paper depends on which of these is chosen, and we carry $\chi \sim 1$ throughout, fixing it only where the seeded computation of Appendix C needs a number.

The fluctuation rent. Equation (9) makes the rate a function of the clock’s own energy. If that energy is not a sharp number but carries a spread σ_E , the rate inherits a spread. Linearizing Eq. (9) in $\bar{E} \rightarrow \bar{E} + \delta E$,

$$\delta\left(\frac{\delta f}{f}\right) = -\frac{G\chi}{Rc^4} \delta E, \quad \left|\delta\left(\frac{\delta f}{f}\right)\right|_{\text{rms}} = \frac{G\chi}{Rc^4} \sigma_E, \quad (10)$$

which is Eq. (3). The mean term of Eq. (9) is a fixed number and is subtracted by calibration against the external standard, exactly as chronometric geodesy subtracts the redshift correction from the known potential; what Eq. (10) adds is the part that fluctuates. It is irreducible only insofar as δE is correlated with the tick: an energy fluctuation statistically independent of the reading would average away over τ like any noise, so the load-bearing claim is that the relevant δE is *conditional on the tick record*. This is where the resource-theoretic identification of the clock as a closed system and the energy-basis character of gravitational decoherence do the work: the internal energy that sets the rate is the same energy whose fluctuation is read out, so the two are perfectly correlated at leading order [5, 18]. We grade Eq. (10) exact given that identification and grade the identification structured, which is why the size of σ_E is bracketed rather than asserted.

The two brackets. The question the brackets answer is: of all the mass-energy in the clock, how much both gravitates and fluctuates conditionally on the tick? Two limits bound the truth.

The *lower bracket* charges only the oscillating degree of freedom. By Proposition 1 the oscillator carries $\sigma_E^{\text{osc}} \geq \hbar \delta\phi/(2\delta t)$, and if only this energy is conditionally correlated with the reading then

$$\mathcal{A} \geq \frac{G\chi}{Rc^4} \sigma_E^{\text{osc}} = \frac{G\hbar\chi}{Rc^4} \cdot \frac{\delta\phi}{2\delta t}, \quad (11)$$

which is Eq. (4). This is the optimistic bound: it assumes the mount, laser, and containment carry energy that either does not fluctuate with the tick or does not gravitate coherently with the oscillator, and it is the smallest defensible rent.

The *upper bracket* charges the clock’s total gravitating mass-energy. If the whole clock of energy $\bar{E}$ shares the tick’s conditional fluctuation — the conservative reading, in which the containment is rigidly correlated with the oscillator it holds — then the relevant spread is the total energy uncertainty $\sigma_E^{\text{tot}} \geq \sigma_E^{\text{osc}}$, and

$$\mathcal{A} \geq \frac{G\chi}{Rc^4} \sigma_E^{\text{tot}} = \frac{G\hbar\chi}{Rc^4} \cdot \frac{\delta\phi}{2\delta t} \cdot \frac{\sigma_E^{\text{tot}}}{\sigma_E^{\text{osc}}}. \quad (12)$$

The two brackets differ only by the ratio $\sigma_E^{\text{tot}}/\sigma_E^{\text{osc}} \geq 1$ of total-to-oscillator energy fluctuation, a number set by engineering, not principle. The true floor lies between Eqs. (11) and (12); both scale as $1/R$ and both carry the

same $\delta\phi/(2\delta t)$, so the bracket shifts the floor's coefficient but never its exponent. That δE decoheres in the energy basis, not the position basis, is what forbids a third, still smaller bracket in which the fluctuation somehow fails to couple to the rate [5, 73], and the G/c^4 weak-coupling scaling of the clock-system interaction is the same one derived independently for interacting clocks [85].

The light-crossing constraint and the optimum.

A single tick is one event only if every part of the clock shares it, which requires the clock's spatial extent to be within the light-crossing distance during the tick, $R \leq c\delta t$ (Eq. (5)); this is the Salecker–Wigner synchronization requirement [79], used here as a constraint on the minimization rather than as a bound in its own right. The rent brackets both decrease monotonically in R , so on the interval $0 < R \leq c\delta t$ the minimum accuracy is attained at the upper endpoint $R^* = c\delta t$. Substituting into Eq. (11),

$$\mathcal{A}_{\min} \geq \frac{G\hbar\chi}{(c\delta t)c^4} \cdot \frac{\delta\phi}{2\delta t} = \frac{G\hbar\chi\delta\phi}{2c^5\delta t^2} = \frac{\chi\delta\phi}{2} \left(\frac{t_P}{\delta t}\right)^2, \quad (13)$$

using $t_P^2 = \hbar G/c^5$, which is Eq. (6). The upper bracket gives the same expression times $\sigma_E^{\text{tot}}/\sigma_E^{\text{osc}}$. The δt that appears is fixed by the demanded phase resolution $\delta\phi$ over the averaging time τ ; for a clock resolving one radian per interrogation, δt is the interrogation time itself, and the floor improves quadratically as that time is lengthened, until the light-crossing size $R^* = c\delta t$ grows to whatever independent bound (laboratory scale, causal horizon) caps it. The exponent in Eq. (13) is fixed; the coefficient carries $\chi \sim 1$, the design choice $\delta\phi$, and the bracket ratio, and the seeded computation of Appendix C fixes all three to land the number quoted in Section 4.

B Speed limits, resource-theory statements, and the foam derivations reconstructed

This appendix collects, with attribution, the inherited results the main text uses, and reconstructs the foam clock bounds so that the accounting objection of Section 6 can be checked against the derivation rather than a paraphrase.

Speed limits. The quantum price rests on the quantum speed limit in its two classical forms. Mandelstam and Tamm bound the evolution time to a distinguishable state by the energy variance, $\tau_{\text{QSL}} = \pi\hbar/(2\sigma_E)$ [49]; Margolus and Levitin bound it instead by the mean energy above the ground state, $\tau_{\text{QSL}} = \pi\hbar/(2E)$, a bound depending only on $\bar{E} - E^0$ and not on the variance [50]. The attainable rate is the tighter,

$\tau_{\text{QSL}} = \max\{\pi\hbar/(2\sigma_E), \pi\hbar/(2E)\}$, and the modern review states both together with the geometric derivations that unify them [23]. The relation used as Eq. (1), $\sigma_E \geq \hbar\delta\phi/(2\delta t)$, is the Mandelstam–Tamm bound rewritten for a phase $\delta\phi$ accrued in time δt . Its correct reading — that δt is an evolution time-scale and not a measurement duration — is forced by Aharonov and Bohm's measurement scheme [3], and the taxonomy of which time-energy relations are legitimate is Busch's [15]. The geometric substrate, that distinguishability is Fisher-information distance on state space, is Braunstein and Caves' [13], and the metrological scalings $1/(\sqrt{N}\tau)$ and $1/(N\tau)$ are the standard-quantum and Heisenberg limits reviewed by Giovannetti, Lloyd, and Maccone [32].

Resource-theory statements. The autonomous-clock resource theory supplies the closed-system definition of the clock the rent derivation needs. Ranković, Liang, and Renner's alternate-ticks game defines a clock's quality operationally, without an external standard [76]; Woods, Silva, and Oppenheim show the finite quasi-ideal clock's back-reaction falls exponentially in its Hilbert-space dimension while its energy grows only linearly [95], and Woods' axiomatic construction subsumes the autonomous and high-accuracy models in one framework [94]. The thermodynamic floor that runs parallel to ours is Erker and collaborators' entropy-per-tick bound [29], sharpened to a universal resolution-accuracy trade-off for memoryless clocks by Meier and collaborators [55] and to a complexity-versus-second-law statement by Schwarzhans and collaborators [83]; Milburn frames the whole subject through phase diffusion in open systems [57], the single-qubit case gives the $k_B T$ -scaled energy-accuracy trade-off [33], and the trade-off has been measured to run within an order of magnitude of its bound [67]. Related dimension-versus-sharpness and imperfect-timekeeping results [56, 86, 96, 97] complete the column. We use these to fix *which* energy is the clock's own; the entropy floor itself is a distinct column and is never added to the gravitational rent.

The foam derivations reconstructed. The Salecker–Wigner argument sets the founding constraint: a clock of mass m used to time an interval T over a distance has a quantum position spread that grows with T unless m is large, so resolving the interval requires a minimum mass [79], the quantum-measurement companion to Wigner's argument that a covariant measurement needs a physical reference obeying its own uncertainty relations [92]. Ng and van Dam combine this with a Schwarzschild-radius ceiling on the mass. A clock of size ℓ must satisfy two requirements: its quantum spread over the light-crossing time ℓ/c is $\delta x_q \gtrsim \sqrt{\hbar\ell/(mc)}$ from Salecker–Wigner, and it must not lie within its own Schwarzschild radius, $\ell \gtrsim Gm/c^2$, i.e. $m \lesssim c^2\ell/G$. Minimizing the total uncertainty over the mass m balances a term falling in m against a

gravitational term rising in m , and the optimum gives

$$\delta\ell \gtrsim (\ell\ell_P^2)^{1/3}, \quad \frac{\delta\ell}{\ell} \gtrsim \left(\frac{\ell_P}{\ell}\right)^{2/3}, \quad (14)$$

the one-third power, read as a clock as $\delta t \gtrsim t_P^{2/3}\tau^{1/3}$ [60, 61]. The holographic motivation attaches through the counting $\tilde{\nu}^2 I^{-1} \lesssim t_P^{-2}$ for the speed $\tilde{\nu}$ and memory I of a computer such as a black hole [59], with Lloyd’s $2\bar{E}/\pi\hbar$ operations per second and the 10^{120} -operation cosmic bound as the same accounting at other scales [46, 47], and the interferometric version through Hogan’s holographic shear noise of amplitude $\sim \sqrt{ct_P L}$ on a region of size L [34, 35]. The step this reconstruction exposes is the one the accounting objects to: the minimized quantity in Eq. (14) is the clock’s *total mass* m , entering through both the Salecker–Wigner spread and the Schwarzschild ceiling. By Section 3 the mass sets the calibrable mean redshift, not the irreducible fluctuation; charging accuracy for it is charging for a systematic that can be subtracted, and the resulting $\tau^{1/3}$ is therefore a bound on an uncalibrated clock, graded structured as an objection. That the enhancement over the bare Planck scale is in any case measurement-scheme-dependent is Baez and Olson’s rod construction, reaching $\delta\ell \geq \ell_P$ with no $\ell^{1/3}$ factor [7]; that a Planck-scale floor follows from quantum mechanics with classical general relativity alone is Calmet, Graesser, and Hsu’s [16]; that the bare Planck length and time are the operational floors is Padmanabhan’s [65]; and that minimal-length scenarios are model-dependent with recurring escape routes is the burden of the Garay and Hossenfelder surveys [31, 37]. The interferometer-as-quantum-gravity-probe proposal that set the program in motion is Amelino-Camelia’s [4]. Our floor, Eq. (13), replaces the minimized total mass with the fluctuation σ_E fixed by the tick’s own resolution, which is why its exponent is the square of $t_P/\delta t$ rather than the one-third power of Eq. (14): the two bounds are computing charges on two different energies, the calibrable mean and the irreducible fluctuation, and only the second survives calibration.

C The seeded computation

The floor’s exponent, the optimal-size locus, and the instrument map come from one deterministic script, `sim_floor.py`, whose source and outputs (`sim_floor_output.txt/.json`) ship with the paper. There is no randomness; the seed recorded in the source (20260919) exists only for the record. The script’s role is to fix the order-unity constants the main text leaves symbolic, so the floor is quoted as a definite power of ten rather than a scaling.

The floor and its exponent

With the Planck time $t_P = \sqrt{\hbar G/c^5} = 5.39 \times 10^{-44}$ s, the floor of Proposition 3 is $\mathcal{A}_{\min}(R) = (\chi \delta\phi/2)(t_P/\delta t)^2$ with $\delta t = R/c$ the light-crossing time and $\chi = \delta\phi = 1$ declared as order unity. The script tabulates it: at a laboratory scale of $R = 1$ m, $\delta t = 3.34 \times 10^{-9}$ s and $\mathcal{A}_{\min} = 1.3 \times 10^{-70}$, which is about 52 orders of magnitude below the demonstrated 9.4×10^{-19} of the Al^+ optical clock. The floor falls as R^{-2} along the optimal-size locus $R^* = c\delta t$, so a larger coherent clock does better, bounded only by the largest feasible apparatus; the absolute laboratory floor is where that locus meets a metre-scale bench. The exponent is the claim and the coefficient is not quoted to better than order unity, which is why the main text and Fig. 2 carry $\mathcal{A}_{\min} \approx 10^{-70}$ rather than a spuriously precise figure. Both brackets of the fluctuation term — the oscillator’s own excitation uncertainty and the clock’s total gravitating energy — are carried in Appendix A; they move the coefficient, not the power of ten.

The map

The instrument table places each flagship clock on the $(R, \text{accuracy})$ plane, from Essen and Parry’s 10^{-10} caesium standard of 1955 through the Cs fountain at 10^{-16} to the Sr and Al^+ optical clocks at $\sim 10^{-19}$ and the projected Th-229 nuclear clock, against the floor curve. The gap between the demonstrated frontier and the floor is about 52 orders of magnitude at laboratory scale: the accuracy the two theories jointly permit and no instrument has used. That gap, drawn in Fig. 3 and closing the Quantum–Relativity Papers, is the quantitative form of the series’ final statement — that the room quantum mechanics and general relativity leave for a clock is vast, and almost entirely unspent.

References

- [1] A. A. Abdo et al. [Fermi GBM and LAT Collaborations] (2009). A Limit on the Variation of the Speed of Light Arising from Quantum Gravity Effects (testing Einstein’s special relativity with GRB 090510). *Nature*, vol. 462, pp. 331–334. <https://arxiv.org/abs/0908.1832>.
- [2] A. Aeppli, K. Kim, W. Warfield, M. S. Safronova, and J. Ye (2024). Clock with 8×10^{-19} Systematic Uncertainty. *Physical Review Letters*, vol. 133, article 023401. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.133.023401>. DOI: 10.1103/PhysRevLett.133.023401.
- [3] Y. Aharonov and D. Bohm (1961). Time in the Quantum Theory and the Uncertainty Relation for Time and Energy. *Physical Review*, vol. 122, p. 1649. <https://journals.aps.org/pr/abstract/>

- 10.1103/PhysRev.122.1649. DOI: 10.1103/PhysRev.122.1649.
- [4] G. Amelino-Camelia (1999). An Interferometric Gravitational Wave Detector as a Quantum-Gravity Apparatus. *Nature*, vol. 398, pp. 216–218. <https://arxiv.org/abs/gr-qc/9808029>.
 - [5] C. Anastopoulos and B. L. Hu (2013). A Master Equation for Gravitational Decoherence: Probing the Textures of Spacetime. *Classical and Quantum Gravity*, vol. 30, article 165007. <https://arxiv.org/abs/1305.5231>.
 - [6] Boulder Atomic Clock Optical Network (BACON) Collaboration: K. Beloy, M. I. Bodine, T. Bothwell, S. M. Brewer, et al. (2021). Frequency Ratio Measurements at 18-Digit Accuracy Using an Optical Clock Network. *Nature*, vol. 591, pp. 564–569. <https://www.nature.com/articles/s41586-021-03253-4>. DOI: 10.1038/s41586-021-03253-4.
 - [7] J. C. Baez and S. J. Olson (2002). Uncertainty in Measurements of Distance. *Classical and Quantum Gravity*, vol. 19, pp. L121–L126. <https://arxiv.org/abs/gr-qc/0201030>.
 - [8] K. Beeks, T. Sikorsky, T. Schumm, J. Thielking, M. V. Okhapkin, and E. Peik (2021). The Thorium-229 Low-Energy Isomer and the Nuclear Clock. *Nature Reviews Physics*, vol. 3, pp. 238–248. <https://www.nature.com/articles/s42254-021-00286-6>. DOI: 10.1038/s42254-021-00286-6.
 - [9] K. Beeks, G. A. Kazakov, F. Schaden, et al. (2024). Fine-Structure Constant Sensitivity of the Th-229 Nuclear Clock Transition. *Nature Communications*, vol. 16, article 9147 (2025). <https://arxiv.org/abs/2407.17300>.
 - [10] S. Bize (2019). The Unit of Time: Present and Future Directions. *Comptes Rendus Physique*, vol. 20, pp. 153–168. <https://doi.org/10.1016/j.crhy.2019.02.002>. DOI: 10.1016/j.crhy.2019.02.002.
 - [11] T. Bothwell, C. J. Kennedy, A. Aepli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye (2022). Resolving the Gravitational Redshift across a Millimetre-Scale Atomic Sample. *Nature*, vol. 602, pp. 420–424. <https://www.nature.com/articles/s41586-021-04349-7>. DOI: 10.1038/s41586-021-04349-7.
 - [12] T. Bothwell, W. Brand, R. Fasano, et al. (2024). Deployment of a Transportable Optical Lattice Clock. <https://arxiv.org/abs/2409.16264>.
 - [13] S. L. Braunstein and C. M. Caves (1994). Statistical Distance and the Geometry of Quantum States. *Physical Review Letters*, vol. 72, p. 3439. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.72.3439>. DOI: 10.1103/PhysRevLett.72.3439.
 - [14] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibbrandt (2019). $^{27}\text{Al}^+$ Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18} . *Physical Review Letters*, vol. 123, article 033201. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.123.033201>. DOI: 10.1103/PhysRevLett.123.033201.
 - [15] P. Busch (2008). The Time–Energy Uncertainty Relation. Chapter 3 in *Time in Quantum Mechanics* (2nd ed.), Lecture Notes in Physics 734, Springer, pp. 73–105. <https://arxiv.org/abs/quant-ph/0105049>. DOI: 10.1007/978-3-540-73473-4_3.
 - [16] X. Calmet, M. Graesser, and S. D. H. Hsu (2004). Minimum Length from Quantum Mechanics and Classical General Relativity. *Physical Review Letters*, vol. 93, article 211101. <https://arxiv.org/abs/hep-th/0405033>.
 - [17] C. J. Campbell, A. G. Radnaev, A. Kuzmich, V. A. Dzuba, V. V. Flambaum, and A. Derevianko (2012). A Single-Ion Nuclear Clock for Metrology at the 19th Decimal Place. *Physical Review Letters*, vol. 108, article 120802. <https://arxiv.org/abs/1110.2490>.
 - [18] E. Castro-Ruiz, F. Giacomini, and Č. Brukner (2017). Entanglement of Quantum Clocks through Gravity. *Proceedings of the National Academy of Sciences*, vol. 114, no. 12, pp. E2303–E2309. <https://arxiv.org/abs/1507.01955>.
 - [19] A. Chou et al. [Holometer Collaboration] (2017). Interferometric Constraints on Quantum Geometrical Shear Noise Correlations. *Classical and Quantum Gravity*, vol. 34, no. 16, article 165005. <https://arxiv.org/abs/1703.08503>.
 - [20] B. Corbeel (2026). Every Interferometer Is Secretly a Clock Comparison. *The Quantum–Relativity Papers, No. 1*, Absolute Digital Publishers. <https://absolutedigitalpublishers.com/papers/quantum-relativity-01-clock-comparison.pdf>.
 - [21] B. Corbeel (2026). An Atom Held in Two Places Develops a Linewidth. *The Quantum–Relativity Papers, No. 2*, Absolute Digital Publishers. <https://absolutedigitalpublishers.com/papers/>

- quantum-relativity-02-gravitational-linewidth.pdf.
- [22] B. Corbeel (2026). The Coldest Places in the Universe Are Auditing Wavefunction Collapse. *The Quantum-Relativity Papers*, No. 5, Absolute Digital Publishers. <https://absolutedigitalpublishers.com/papers/quantum-relativity-05-collapse-heating-ledger.pdf>.
- [23] S. Deffner and S. Campbell (2017). Quantum Speed Limits: From Heisenberg’s Uncertainty Principle to Optimal Quantum Control. *Journal of Physics A: Mathematical and Theoretical*, vol. 50, article 453001. <https://arxiv.org/abs/1705.08023>.
- [24] P. Delva, H. Denker, and G. Lion (2018). Chronometric Geodesy: Methods and Applications. In D. Puetzfeld et al. (Eds.), *Relativistic Geodesy: Foundations and Applications* (2019), pp. 25–85. <https://arxiv.org/abs/1804.09506>.
- [25] N. Dimarcq, M. Gertsvolf, G. Miletì, S. Bize, C. W. Oates, E. Peik, et al. (2024). Roadmap towards the Redefinition of the Second. *Metrologia*, vol. 61, article 012001. <https://doi.org/10.1088/1681-7575/ad17d2>. DOI: 10.1088/1681-7575/ad17d2.
- [26] L. Diósi (1984). Gravitation and Quantum-Mechanical Localization of Macro-Objects. *Physics Letters A*, vol. 105, pp. 199–202. <https://arxiv.org/abs/1412.0201>. DOI: 10.1016/0375-9601(84)90397-9.
- [27] L. Diósi (1989). Models for Universal Reduction of Macroscopic Quantum Fluctuations. *Physical Review A*, vol. 40, pp. 1165–1174. <https://journals.aps.org/pr/abstract/10.1103/PhysRevA.40.1165>. DOI: 10.1103/PhysRevA.40.1165.
- [28] R. Elwell, C. Schneider, J. Jeet, J. E. S. Terhune, H. W. T. Morgan, A. N. Alexandrova, H. B. Tran Tan, A. Derevianko, and E. R. Hudson (2024). Laser Excitation of the ^{229}Th Nuclear Isomeric Transition in a Solid-State Host. *Physical Review Letters*, vol. 133, article 013201. <https://arxiv.org/abs/2404.12311>.
- [29] P. Erker, M. T. Mitchison, R. Silva, M. P. Woods, N. Brunner, and M. Huber (2017). Autonomous Quantum Clocks: Does Thermodynamics Limit Our Ability to Measure Time? *Physical Review X*, vol. 7, article 031022. <https://arxiv.org/abs/1609.06704>. DOI: 10.1103/PhysRevX.7.031022.
- [30] L. Essen and J. V. L. Parry (1955). An Atomic Standard of Frequency and Time Interval: A Cæsium Resonator. *Nature*, vol. 176, pp. 280–282. <https://www.nature.com/articles/176280a0>. DOI: 10.1038/176280a0.
- [31] L. J. Garay (1995). Quantum Gravity and Minimum Length. *International Journal of Modern Physics A*, vol. 10, pp. 145–166. <https://arxiv.org/abs/gr-qc/9403008>.
- [32] V. Giovannetti, S. Lloyd, and L. Maccone (2011). Advances in Quantum Metrology. *Nature Photonics*, vol. 5, pp. 222–229. <https://arxiv.org/abs/1102.2318>. DOI: 10.1038/nphoton.2011.35.
- [33] M. Gopalkrishnan, V. Kandula, P. Sriram, A. Deshpande, and B. Muralidharan (2017). A Bayesian View of Single-Qubit Clocks, and an Energy versus Accuracy Tradeoff. *Physical Review A*, vol. 96, article 032339. <https://arxiv.org/abs/1602.00508>. DOI: 10.1103/PhysRevA.96.032339.
- [34] C. J. Hogan (2008). Measurement of Quantum Fluctuations in Geometry. *Physical Review D*, vol. 77, article 104031. <https://arxiv.org/abs/0712.3419>.
- [35] C. J. Hogan (2012). Interferometers as Probes of Planckian Quantum Geometry. *Physical Review D*, vol. 85, article 064007. <https://arxiv.org/abs/1002.4880>.
- [36] M. A. Hohensee, S. Chu, A. Peters, and H. Müller (2012). Comment on “Does an Atom Interferometer Test the Gravitational Redshift at the Compton Frequency?”. *Classical and Quantum Gravity*, vol. 29, article 048001. <https://arxiv.org/abs/1112.6039>.
- [37] S. Hossenfelder (2013). Minimal Length Scale Scenarios for Quantum Gravity. *Living Reviews in Relativity*, vol. 16, article 2. <https://arxiv.org/abs/1203.6191>.
- [38] N. Huntemann, C. Sanner, B. Lipphardt, Chr. Tamm, and E. Peik (2016). Single-Ion Atomic Clock with 3×10^{-18} Systematic Uncertainty. *Physical Review Letters*, vol. 116, article 063001. <https://arxiv.org/abs/1602.03908>. DOI: 10.1103/PhysRevLett.116.063001.
- [39] S. Khandelwal, M. P. E. Lock, and M. P. Woods (2020). Universal Quantum Modifications to General Relativistic Time Dilation in Delocalised Clocks. *Quantum*, vol. 4, article 309. <https://arxiv.org/abs/1904.02178>.
- [40] S. B. Koller, J. Grotti, St. Vogt, A. Al-Masoudi, S. Dörscher, S. Häfner, U. Sterr, and Ch. Lisdat (2017). Transportable Optical Lattice

- Clock with 7×10^{-17} Uncertainty. *Physical Review Letters*, vol. 118, article 073601. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.118.073601>. DOI: 10.1103/PhysRevLett.118.073601.
- [41] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin (2014). A Quantum Network of Clocks. *Nature Physics*, vol. 10, pp. 582–587. <https://arxiv.org/abs/1310.6045>.
- [42] S. Kraemer, J. Moens, M. Athanasakis-Kaklamanakis, et al. (2023). Observation of the Radiative Decay of the ^{229}Th Nuclear Clock Isomer. *Nature*, vol. 617, pp. 706–710. <https://arxiv.org/abs/2209.10276>. DOI: 10.1038/s41586-023-05894-z.
- [43] O. Kwon and C. J. Hogan (2016). Interferometric Tests of Planckian Quantum Geometry Models. *Classical and Quantum Gravity*, vol. 33, no. 10, article 105004. <https://arxiv.org/abs/1410.8197>.
- [44] R. Lieu and L. W. Hillman (2003). The Phase Coherence of Light from Extragalactic Sources: Direct Evidence against First-Order Planck-Scale Fluctuations in Time and Space. *Astrophysical Journal Letters*, vol. 585, pp. L77–L80. <https://arxiv.org/abs/astro-ph/0301184>.
- [45] C. Lisdat, G. Grosche, N. Quintin, et al. (2016). A Clock Network for Geodesy and Fundamental Science. *Nature Communications*, vol. 7, article 12443. <https://www.nature.com/articles/ncomms12443>. DOI: 10.1038/ncomms12443.
- [46] S. Lloyd (2000). Ultimate Physical Limits to Computation. *Nature*, vol. 406, pp. 1047–1054. <https://arxiv.org/abs/quant-ph/9908043>.
- [47] S. Lloyd (2002). Computational Capacity of the Universe. *Physical Review Letters*, vol. 88, article 237901. <https://arxiv.org/abs/quant-ph/0110141>.
- [48] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt (2015). Optical Atomic Clocks. *Reviews of Modern Physics*, vol. 87, pp. 637–701. <https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.87.637>. DOI: 10.1103/RevModPhys.87.637.
- [49] L. Mandelstam and I. Tamm (1945). The Uncertainty Relation between Energy and Time in Non-Relativistic Quantum Mechanics. *Journal of Physics (USSR)*, vol. 9, pp. 249–254.
- [50] N. Margolus and L. B. Levitin (1998). The Maximum Speed of Dynamical Evolution. *Physica D*, vol. 120, pp. 188–195. <https://arxiv.org/abs/quant-ph/9710043>. DOI: 10.1016/S0167-2789(98)00054-2.
- [51] M. C. Marshall, D. A. Rodriguez Castillo, W. J. Arthur-Dworschack, et al. (2025). High-Stability Single-Ion Clock with 5.5×10^{-19} Systematic Uncertainty. *Physical Review Letters*, vol. 135, article 023401. <https://journals.aps.org/prl/abstract/10.1103/hb3c-dk28>. DOI: 10.1103/hb3c-dk28.
- [52] G. D. Martinez, C. Li, A. Staron, J. Kitching, C. Raman, and W. R. McGehee (2023). A Chip-Scale Atomic Beam Clock. <https://arxiv.org/abs/2303.11458>.
- [53] W. F. McGrew, X. Zhang, R. J. Fasano, et al. (2018). Atomic Clock Performance Enabling Geodesy below the Centimetre Level. *Nature*, vol. 564, pp. 87–90. <https://arxiv.org/abs/1807.11282>. DOI: 10.1038/s41586-018-0738-2.
- [54] T. E. Mehlstäubler, G. Grosche, C. Lisdat, P. O. Schmidt, and H. Denker (2018). Atomic Clocks for Geodesy. *Reports on Progress in Physics*, vol. 81, article 064401. <https://arxiv.org/abs/1803.01585>.
- [55] F. Meier, E. Schwarzhans, P. Erker, and M. Huber (2023). Fundamental Accuracy-Resolution Trade-off for Timekeeping Devices. *Physical Review Letters*, vol. 131, article 220201. <https://arxiv.org/abs/2301.05173>. DOI: 10.1103/PhysRevLett.131.220201.
- [56] F. Meier, M. Huber, P. Erker, and J. Xuereb (2024). Autonomous Quantum Processing Unit: An Autonomous Thermal Computing Machine & Its Physical Limitations. *Reports on Progress in Physics*, vol. 89, article 7 (2026). <https://arxiv.org/abs/2402.00111>. DOI: 10.1088/1361-6633/ae7b82.
- [57] G. J. Milburn (2020). The Thermodynamics of Clocks. *Contemporary Physics*, vol. 61, no. 2, pp. 69–95. <https://arxiv.org/abs/2007.02217>. DOI: 10.1080/00107514.2020.1837471.
- [58] H. Müller, A. Peters, and S. Chu (2010). A Precision Measurement of the Gravitational Redshift by the Interference of Matter Waves. *Nature*, vol. 463, pp. 926–929. <https://doi.org/10.1038/nature08776>. DOI: 10.1038/nature08776.
- [59] Y. J. Ng (2001). From Computation to Black Holes and Space-Time Foam. *Physical Review Letters*, vol. 86, p. 2946 (Erratum: vol. 88, article 139902, 2002). <https://arxiv.org/abs/gr-qc/0006105>.

- [60] Y. J. Ng and H. van Dam (1994). Limit to Space-Time Measurement. *Modern Physics Letters A*, vol. 9, no. 4, pp. 335–340. <https://doi.org/10.1142/S0217732394000356>. DOI: 10.1142/S0217732394000356.
- [61] Y. J. Ng and H. van Dam (2000). Measuring the Foaminess of Space-Time with Gravity-Wave Interferometers. *Foundations of Physics*, vol. 30, no. 6, pp. 795–805. <https://arxiv.org/abs/gr-qc/9906003>.
- [62] I. Nosske, C. Vishwakarma, T. Lücke, J. Rahm, N. Poudel, S. Weyers, E. Benkler, S. Dörscher, and C. Lisdat (2025). Transportable Strontium Lattice Clock with 4×10^{-19} Blackbody Radiation Shift Uncertainty. <https://arxiv.org/abs/2507.14030>.
- [63] T. Ooi, J. F. Doyle, C. Zhang, J. S. Higgins, J. Ye, K. Beeks, T. Sikorsky, and T. Schumm (2025). Frequency Reproducibility of Solid-State Th-229 Nuclear Clocks. *Nature*, vol. 650, pp. 72–78 (2026). <https://arxiv.org/abs/2507.01180>. DOI: 10.1038/s41586-025-09999-5.
- [64] J. Paczos, K. Dębski, P. T. Grochowski, A. R. H. Smith, and A. Dragan (2024). Quantum Time Dilation in a Gravitational Field. *Quantum*, vol. 8, article 1338. <https://arxiv.org/abs/2204.10609>.
- [65] T. Padmanabhan (1987). Limitations on the Operational Definition of Space-Time Events and Quantum Gravity. *Classical and Quantum Gravity*, vol. 4, no. 4, pp. L107–L113. <https://doi.org/10.1088/0264-9381/4/4/007>. DOI: 10.1088/0264-9381/4/4/007.
- [66] S. Park and J. H. Rhee (2024). CSAC Drift Modeling Considering GPS Signal Quality in the Case of GPS Signal Unavailability. <https://arxiv.org/abs/2409.00594>.
- [67] A. N. Pearson, Y. Guryanova, P. Erker, E. A. Laird, G. A. D. Briggs, M. Huber, and N. Ares (2021). Measuring the Thermodynamic Cost of Timekeeping. *Physical Review X*, vol. 11, article 021029. <https://arxiv.org/abs/2006.08670>. DOI: 10.1103/PhysRevX.11.021029.
- [68] E. Pedrozo-Peñafiel, S. Colombo, C. Shu, et al. (2020). Entanglement on an Optical Atomic-Clock Transition. *Nature*, vol. 588, pp. 414–418. <https://arxiv.org/abs/2006.07501>.
- [69] E. Peik, T. Schumm, M. S. Safronova, A. Pálffy, J. Weitenberg, and P. G. Thirolf (2021). Nuclear Clocks for Testing Fundamental Physics. *Quantum Science and Technology*, vol. 6, article 034002. <https://arxiv.org/abs/2012.09304>. DOI: 10.1088/2058-9565/abe9c2.
- [70] E. Peik and Chr. Tamm (2003). Nuclear Laser Spectroscopy of the 3.5 eV Transition in Th-229. *Europhysics Letters*, vol. 61, no. 2, pp. 181–186. <https://doi.org/10.1209/epl/i2003-00210-x>. DOI: 10.1209/epl/i2003-00210-x.
- [71] R. Penrose (1996). On Gravity’s Role in Quantum State Reduction. *General Relativity and Gravitation*, vol. 28, no. 5, pp. 581–600. <http://link.springer.com/10.1007/BF02105068>. DOI: 10.1007/BF02105068.
- [72] E. S. Perlman, S. A. Rappaport, W. A. Christiansen, Y. J. Ng, J. DeVore, and D. Pooley (2015). New Constraints on Quantum Gravity from X-Ray and Gamma-Ray Observations. *The Astrophysical Journal*, vol. 805, no. 1, article 10. <https://arxiv.org/abs/1411.7262>.
- [73] I. Pikovski, M. Zych, F. Costa, and Č. Brukner (2015). Universal Decoherence Due to Gravitational Time Dilation. *Nature Physics*, vol. 11, pp. 668–672. <https://arxiv.org/abs/1311.1095>.
- [74] R. Ragazzoni, M. Turatto, and W. Gaessler (2003). Lack of Observational Evidence for Quantum Structure of Space-Time at Planck Scales. *The Astrophysical Journal*, vol. 587, pp. L1–L4. <https://arxiv.org/abs/astro-ph/0303043>.
- [75] N. F. Ramsey (1950). A Molecular Beam Resonance Method with Separated Oscillating Fields. *Physical Review*, vol. 78, pp. 695–699. <https://journals.aps.org/pr/abstract/10.1103/PhysRev.78.695>. DOI: 10.1103/PhysRev.78.695.
- [76] S. Ranković, Y.-C. Liang, and R. Renner (2015). Quantum Clocks and Their Synchronisation — The Alternate Ticks Game. <https://arxiv.org/abs/1506.01373>.
- [77] J. D. Roslund, A. Cingöz, W. D. Lunden, et al. (2023). Optical Clocks at Sea. <https://arxiv.org/abs/2308.12457>.
- [78] A. Roura (2020). Gravitational Redshift in Quantum-Clock Interferometry. *Physical Review X*, vol. 10, no. 2, article 021014. <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.10.021014>. DOI: 10.1103/PhysRevX.10.021014.
- [79] H. Salecker and E. P. Wigner (1958). Quantum Limitations of the Measurement of Space-Time Distances. *Physical Review*, vol. 109, no. 2, pp. 571–577. <https://doi.org/10.1103/PhysRev.109.571>. DOI: 10.1103/PhysRev.109.571.

- [80] C. Sanner, N. Huntemann, R. Lange, Chr. Tamm, E. Peik, M. S. Safronova, and S. G. Porsev (2019). Optical Clock Comparison for Lorentz Symmetry Testing. *Nature*, vol. 567, pp. 204–208. <https://www.nature.com/articles/s41586-019-0972-2>. DOI: 10.1038/s41586-019-0972-2.
- [81] M. Schioppo, J. Kronjäger, A. Silva, et al. (2022). Comparing Ultrastable Lasers at 7×10^{-17} Fractional Frequency Instability through a 2220 km Optical Fibre Network. *Nature Communications*, vol. 13, article 212. <https://www.nature.com/articles/s41467-021-27884-3>. DOI: 10.1038/s41467-021-27884-3.
- [82] M. Schulte, C. Lisdat, P. O. Schmidt, U. Sterr, and K. Hammerer (2020). Optimal Clock Stability with Spin-Squeezed States and Brownian Frequency Noise. *Nature Communications*, vol. 11, article 5955. <https://arxiv.org/abs/1911.00882>.
- [83] E. Schwarzhans, M. P. E. Lock, P. Erker, N. Friis, and M. Huber (2021). Autonomous Temporal Probability Concentration: Clockworks and the Second Law of Thermodynamics. *Physical Review X*, vol. 11, article 011046. <https://arxiv.org/abs/2007.01307>. DOI: 10.1103/PhysRevX.11.011046.
- [84] B. Seiferle, L. von der Wense, P. V. Bilous, et al. (2019). Energy of the 229Th Nuclear Clock Transition. *Nature*, vol. 573, pp. 243–246. <https://arxiv.org/abs/1905.06308>. DOI: 10.1038/s41586-019-1533-4.
- [85] A. R. H. Smith and M. Ahmadi (2019). Quantizing Time: Interacting Clocks and Systems. *Quantum*, vol. 3, article 160. <https://arxiv.org/abs/1712.00081>.
- [86] S. Strelchuk and M. P. Woods (2021). Measuring Time with Stationary Quantum Clocks. <https://arxiv.org/abs/2106.07684>.
- [87] P. G. Thirolf, B. Seiferle, and L. von der Wense (2019). The 229-Thorium Isomer: Doorway to the Road from the Atomic Clock to the Nuclear Clock. *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 52, article 203001. <https://doi.org/10.1088/1361-6455/ab29b8>. DOI: 10.1088/1361-6455/ab29b8.
- [88] J. Tiedau, M. V. Okhapkin, K. Zhang, et al. (2024). Laser Excitation of the Th-229 Nucleus. *Physical Review Letters*, vol. 132, article 182501. <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.132.182501>.
- [89] L. Toscani De Col, T. Riebner, I. Morawetz, et al. (2026). A Thorium-229 Optical Nuclear Clock with Feedback Loop. <https://arxiv.org/abs/2606.04997>.
- [90] I. Ushijima, M. Takamoto, M. Das, T. Ohkubo, and H. Katori (2015). Cryogenic Optical Lattice Clocks. *Nature Photonics*, vol. 9, pp. 185–189. <https://www.nature.com/articles/nphoton.2015.5>. DOI: 10.1038/nphoton.2015.5.
- [91] L. von der Wense, B. Seiferle, M. Laatiaoui, et al. (2016). Direct Detection of the 229Th Nuclear Clock Transition. *Nature*, vol. 533, pp. 47–51. <https://arxiv.org/abs/1710.11398>.
- [92] E. P. Wigner (1957). Relativistic Invariance and Quantum Phenomena. *Reviews of Modern Physics*, vol. 29, no. 3, pp. 255–268. <https://doi.org/10.1103/RevModPhys.29.255>. DOI: 10.1103/RevModPhys.29.255.
- [93] P. Wolf, L. Blanchet, C. J. Bordé, S. Reynaud, C. Salomon, and C. Cohen-Tannoudji (2011). Does an Atom Interferometer Test the Gravitational Redshift at the Compton Frequency? *Classical and Quantum Gravity*, vol. 28, article 145017. <https://arxiv.org/abs/1012.1194>.
- [94] M. P. Woods (2021). Autonomous Ticking Clocks from Axiomatic Principles. *Quantum*, vol. 5, article 381. <https://arxiv.org/abs/2005.04628>. DOI: 10.22331/q-2021-01-17-381.
- [95] M. P. Woods, R. Silva, and J. Oppenheim (2019). Autonomous Quantum Machines and the Finite-Sized Quasi-Ideal Clock. *Annales Henri Poincaré* (2018). <https://arxiv.org/abs/1607.04591>. DOI: 10.1007/s00023-018-0736-9.
- [96] J. Xuereb, F. Meier, P. Erker, M. T. Mitchison, and M. Huber (2023). The Impact of Imperfect Timekeeping on Quantum Control. *Physical Review Letters*, vol. 131, article 160204. <https://arxiv.org/abs/2301.10767>. DOI: 10.1103/PhysRevLett.131.160204.
- [97] Y. Yang and R. Renner (2020). Ultimate Limit on Time Signal Generation. <https://arxiv.org/abs/2004.07857>.
- [98] C. Zhang, T. Ooi, J. S. Higgins, et al. (2024). Frequency Ratio of the 229mTh Nuclear Isomeric Transition and the 87Sr Atomic Clock. *Nature*, vol. 633, pp. 63–70. <https://arxiv.org/abs/2406.18719>.
- [99] X. Zheng, J. Dolde, V. Lochab, B. N. Merriam, H. Li, and S. Kolkowitz (2022). Differential Clock Comparisons with a Multiplexed Optical Lattice Clock. *Nature*, vol. 602, pp. 425–430. <https://www.nature.com/articles/s41586-021-04344-y>. DOI: 10.1038/s41586-021-04344-y.

- [100] M. Zych, F. Costa, I. Pikovski, and Č. Brukner (2011). Quantum Interferometric Visibility as a Witness of General Relativistic Proper Time. *Nature Communications*, vol. 2, article 505. <https://arxiv.org/abs/1105.4531>.