It's About Time (not Space!)

Why the Apple Falls, the Starlight Bends, and the Clock Stops at a Horizon

Sridhar Tayur

July 2026

Abstract

Every popular account of general relativity shows the same picture: a bowling ball on a rubber sheet, warping space, with marbles rolling toward it. It is vivid, it is everywhere – and it is a surprisingly poor guide to why an apple falls. This musing derives Newton's law of gravity, $\ddot{\mathbf{x}} = -\nabla\Phi$, from a completely different and far more honest starting point: clocks tick at different rates at different heights, and a free object simply drifts toward wherever time is passing more slowly. No metric tensor, no curvature, no differential geometry – just special relativity's time dilation formula, one Taylor expansion, and the calculus of variations that any second-year undergraduate already owns. What falls out is the oldest law in physics, seen for the first time as a consequence of clocks. As a coda, we then deliberately push past the two places this simplicity has to end. First, light, where the slow-speed approximation fails completely: restoring the spatial part of the metric – the one ingredient this musing otherwise avoids – exactly doubles the deflection of starlight grazing the sun, turning Einstein's own too-small 1911 estimate into the correct 1915 prediction confirmed at the 1919 eclipse. Second, the black hole horizon, where the weak-field *approximation* itself fails: dropping the linearization entirely and using the exact time dilation formula shows a clock's tick rate falling to zero at a finite radius, the same clocks-only logic pushed to its breaking point rather than abandoned. A closing appendix supplies the one piece of mathematical scaffolding assumed throughout – what a metric tensor actually is, built from nothing but tilted, unequal-length coordinate axes. Intended for advanced undergraduates or beginning graduate students comfortable with special relativity and Newtonian mechanics; no prior exposure to general relativity is assumed.

Contents

---

1 The Picture Everyone Is Given, and Why It Is Wrong

Ask anyone who has watched a science documentary how gravity works, and you will get the same answer: mass bends space, and things roll along the resulting slope, the way a marble spirals into a dip in a stretched rubber sheet. The image is unforgettable – and it is not without merit; it is a memorable visualization of spatial curvature, and a fine one for some genuinely relativistic phenomena. But it is a surprisingly poor guide to the specific question of why an apple falls, and worse, it is mildly circular as a demonstration: the rubber sheet only pulls the marble inward because gravity is pulling the marble down onto the sheet in the first place. The demonstration explains gravity by presupposing gravity.

There is a second, quieter problem with the picture, one that has nothing to do with circularity: it draws attention to the *wrong* kind of curvature. For an apple falling from a tree, or a satellite in low orbit, or you standing on a bathroom scale, essentially all of the effect comes not from the curvature of *space* but from the curvature of *time*. In the solar system, spatial curvature is small enough to be safely ignored for this purpose; temporal curvature is not.

This musing is a demonstration of that fact, built from the ground up, using nothing beyond special relativity and ordinary calculus. By the end, Newton's inverse-square law – the law Newton himself simply posited, unexplained, in 1687 – will have been *derived* from a single physical idea: a clock lower in a gravitational well runs slow, and a free body drifts toward wherever its own wristwatch would tick the most.

Where this musing sits. No prior lecture in this series is required. We assume only: (i) the special-relativistic formula for time dilation due to motion, and (ii) first-year single-variable calculus, up through the idea of a Taylor series and the notion of minimizing a quantity by calculus. Section 7 uses the calculus of variations, but we build it from scratch with worked arithmetic rather than assuming it.

2 Two Kinds of Time Dilation, Reviewed From Scratch

2.1 Dilation from Motion

Special relativity (1905) tells us that a moving clock ticks slowly compared to a stationary one. If a clock moves at constant speed v, and an interval dt elapses on the clock of a stationary bookkeeper, the moving clock itself only advances by

$$d\tau = dt\sqrt{1-\frac{v^2}{c^2}}\,, \tag{1}$$

where τ ("proper time") is what the moving clock itself reads, and $c \approx 3\times10^8$ m/s is the speed of light. For $v \ll c$, a first-order Taylor expansion of the square root gives the more usable approximate form

$$d\tau \approx dt\left(1-\frac{v^2}{2c^2}\right). \tag{2}$$

Numerical check. A commercial jet cruises at $v \approx 250$ m/s. Then $v^2/2c^2 \approx (250)^2/(2\times(3\times10^8)^2) \approx 3.5\times10^{-14}$. Over a 10-hour flight (3.6×10^4 s), the pilot's watch falls behind a ground clock by about $3.6\times10^4\times3.5\times10^{-14} \approx 1.3\times10^{-9}$ s – about 1.3 nanoseconds. Tiny, but this exact effect was measured directly in the 1971 Hafele–Keating experiment, using cesium atomic clocks flown around the world on commercial airliners [6, 7].

2.2 Dilation from Gravity: Einstein's 1907 Argument

Equation (1) is pure special relativity – flat spacetime, no gravity. The genuinely new idea, which Einstein had already by 1907 (eight years before the full field equations of general relativity), is that *gravity itself* dilates time, and you can see this using nothing but special relativity and one physical postulate: the **equivalence principle** [1].

The equivalence principle says: being at rest in a gravitational field of strength g is locally indistinguishable from being inside a rocket accelerating at g through empty space, far from any mass. Einstein called this "the happiest thought of my life," because it lets you analyze gravity using the machinery of special relativity applied to accelerated frames – no new theory needed, yet.

Consider a rocket of height h, accelerating upward at rate g. A photon is emitted from the floor and travels to the ceiling.

Derivation: gravitational redshift from an accelerating rocket.

Let the photon leave the floor at $t = 0$, when the floor (and the whole rocket) is momentarily at rest in some external inertial frame. The photon takes a time $t \approx h/c$ to reach the ceiling. During that brief interval, the rocket – and the ceiling detector – has picked up a velocity

$$v = g \cdot t \approx \frac{gh}{c}\,.$$

The ceiling detector is therefore now receding from the photon's original launch frame at speed v, and a receding observer sees a redshifted (lower) frequency, by the standard (non-relativistic, low-v) Doppler formula:

$$\frac{\Delta f}{f} \approx -\frac{v}{c} = -\frac{gh}{c^2}\,.$$

Now translate this back into the language of gravity, using the equivalence principle: an accelerating rocket floor *is* a gravitational field of strength g, and gh is simply the difference in gravitational potential Φ between floor and ceiling, i.e. $gh = \Delta\Phi$. So

$$\frac{\Delta f}{f} = -\frac{\Delta\Phi}{c^2}\,. \tag{3}$$

A photon climbing out of a gravitational well loses frequency – it redshifts. Equivalently, since frequency is a clock (each tick is one oscillation), a clock sitting deeper in a potential well (more negative Φ) ticks *slower* than a clock higher up. This is gravitational time dilation, and note carefully: we derived it using only flat-spacetime special relativity plus one postulate about accelerated frames. No curved spacetime has been invoked.

Equation (3), restated as a statement about elapsed proper time rather than frequency, says that a clock sitting at potential Φ (relative to some reference, with $\Phi < 0$ near a mass, $\Phi \to 0$ far away) runs at the rate

$$d\tau \approx dt\left(1 + \frac{\Phi}{c^2}\right). \tag{4}$$

Sanity check on the sign. Near Earth, $\Phi = -GM/r$ is negative, and gets more negative the closer you are to Earth's center. Equation (4) then says $d\tau/dt < 1$ close to Earth – clocks run slow near the ground compared to clocks far away. This is exactly right, and it is not a small effect for engineering: GPS satellite clocks, sitting about 20,000 km up where Φ is less negative, run *faster* than ground clocks by about 45 microseconds/day from this effect alone (partially offset by a slower-running correction from their orbital speed, via Eq. (2), of about -7 microseconds/day). If this net $+38$ microseconds/day drift were not corrected for in firmware, GPS position errors would accumulate at roughly 10 km per day [8].

3 One Clock, Two Effects: Combining Them

A real free-falling object is simultaneously moving (special-relativistic dilation, Eq. (2)) and sitting at some gravitational potential (gravitational dilation, Eq. (4)). To leading order, for weak fields and slow speeds, these two small corrections simply add (their product, along with the terms dropped in each individual Taylor expansion, is of order $(\Phi/c^2)^2$, $(v/c)^4$, and $(\Phi/c^2)(v/c)^2$ – for the solar system, $\Phi/c^2 \sim 10^{-8}$ and $(v/c)^2 \lesssim 10^{-8}$, so these discarded terms are $\lesssim 10^{-16}$, some sixteen orders of magnitude smaller than the leading corrections we keep, utterly negligible for planetary fields and non-relativistic velocities):

$$\boxed{d\tau \approx dt\left(1+\frac{\Phi}{c^2}-\frac{v^2}{2c^2}\right)} \tag{5}$$

This single line is the entire physical content of the musing. Everything from here on is algebra and calculus applied to Eq. (5). It is worth pausing on how little was needed to get here: one experimental fact about moving clocks (1905 special relativity) and one thought experiment about accelerating rockets (1907 equivalence principle). Full general relativity – curved spacetime, the metric tensor $g_{\mu\nu}$, the Einstein field equations – is nowhere in sight, and will not be needed.

4 The Principle of Maximal Aging

Nature, it turns out, has a habit of organizing itself around a simple question: out of all the ways something *could* move between two points, which one does it actually take? Fermat answered this for light in 1662: it takes the path of least time. Maupertuis, Euler, and Lagrange answered it for mechanics in the 1740s: a moving body takes the path that extremizes a quantity later called the action. Hamilton tightened the whole picture in the 1830s into a single variational statement broad enough to cover both. Einstein's contribution, a century later, was to ask the same kind of question about gravity – and the answer, once again, turns out to be a statement about time.

Here is the one genuinely new postulate of this musing, due in modern pedagogical form to Taylor and Wheeler [5], though the physical content goes back to Einstein: *a free particle – one acted on by gravity alone, no rockets, no rope, no engine – follows the path between two events that makes the elapsed proper time τ on its own wristwatch as large as possible*, compared to every nearby alternative path connecting the same two events.

For the reader who already knows some GR. Maximal aging is not a free-floating axiom. In the full theory, a free particle moves on a geodesic – the curve extremizing $\tau = \int\sqrt{-g_{\mu\nu}\,dx^\mu dx^\nu}$ between two fixed events, with τ coming out *maximal* because any nearby path that wiggles through space always logs *less* proper time (the twin-paradox inequality, in different clothes). Equation (5) is just that same expression, written out to leading order for a slow, weak-field observer. So everything below is ordinary geodesic motion, seen in its simplest limit – not a separate or ad hoc idea.

This is a strange-sounding rule at first, so it is worth building intuition with pure arithmetic before any calculus at all.

Arithmetic example: dropping a ball for one second.

Suppose you release a ball from height h_0 above the ground, and it returns to your hand exactly $T = 1$ second later (idealize: no air resistance; imagine it bounces elastically off a partner's hand at the bottom and returns, or just consider the up-going half separately – the arithmetic point survives). Compare two candidate paths a fictitious particle could take between the same start and end events (same starting height, same 1-second total elapsed time on an external clock):

Path A (the real one): free fall. The ball leaves your hand, rises, and returns, spending equal time near the top (small v, higher Φ) and near the bottom (large v, lower Φ).

Path B (a fictitious alternative): constant velocity. A particle that moves through the same total displacement at perfectly constant speed the whole second, never slowing at the top.

Equation (5) says elapsed proper time is $\tau = \int dt\,(1 + \Phi/c^2 - v^2/2c^2)$. Near Earth's surface both Φ/c^2 and $v^2/2c^2$ are of order 10^{-16} per unit height/speed relevant here, so the *numbers* are unmeasurably small in this everyday example – but the *sign structure* is exactly what matters, and it is the same sign structure at any scale. Path A, real free fall, spends more of its time at small v (near the top of the arc, where the ball momentarily stops) than Path B does, and $-v^2/2c^2$ penalizes proper time less when v is small. Free fall is precisely the trajectory that optimally trades off "spend time up high, where Φ is more favorable and clocks run faster" against "don't dawdle so long that you pick up penalizing speed getting back down." The parabola is not an arbitrary shape – it is the unique shape that maximizes $\int(1 + \Phi/c^2 - v^2/2c^2)\,dt$.

The content of maximal aging, stated as an extremization problem, is:

$$\tau = \int dt + \frac{1}{c^2}\int\left(\Phi - \frac{v^2}{2}\right)dt \quad\longrightarrow\quad \text{maximize over all paths.} \tag{6}$$

The first term, $\int dt$, is the same fixed number (T, the total external-clock time between start and end events) for every candidate path – it does not distinguish paths at all. So maximizing τ is equivalent to maximizing the second term alone:

$$\text{maximize}\int\left(\Phi - \frac{v^2}{2}\right)dt \iff \text{minimize}\int\left(\frac{v^2}{2} - \Phi\right)dt\,. \tag{7}$$

5 Recognizing an Old Friend

Equation (7) should look familiar to anyone who has taken a mechanics course, even without relativity in sight. The quantity Φ appearing throughout this musing is already, by construction (Section 2.2, via the equivalence principle), the ordinary Newtonian gravitational potential: negative near a mass, $\Phi = -GM/r$ for a point mass, increasing toward zero far away. So define, per unit mass,

$$T \equiv \frac{v^2}{2} \quad \text{(kinetic energy per unit mass)}, \qquad U \equiv \Phi \quad \text{(potential energy per unit mass, so the usual potential ener}$$

so that the quantity being minimized in Eq. (7) is exactly

$$S = \int (T - U)\,dt = \int L\,dt\,, \qquad L \equiv T - U = \frac{v^2}{2} - \Phi\,, \tag{8}$$

which is precisely the classical Lagrangian and the classical action integral, first written down by Euler, Lagrange, and Maupertuis in the 1740s [9, 10] – more than 150 years before special or general relativity existed. We have arrived at the principle of least action not as an axiom handed down from above, but as a direct, derived consequence of clocks running at different rates. What generations of students are simply told to accept ("postulate: nature minimizes the action") turns out, at least in the weak-field slow-speed limit, to be a restatement of maximal aging on a wristwatch.

The conceptual climax of this musing, in one sentence. Least action is not an independent principle of nature, standing beside relativity as a second, unrelated axiom – it is the weak-field shadow of maximal aging.

A note on "least" versus "most." Students sometimes find it odd that Section 4 says *maximize* proper time while Section 5 says *minimize* the action. There is no contradiction – maximizing τ in Eq. (6) is algebraically identical to minimizing $S = \int (T - U)\, dt$ in Eq. (7), because of the overall minus sign introduced when we factored out the fixed term $\int dt$. Both statements describe the same underlying trajectory.

6 Turning the Crank: Euler–Lagrange

To find the path $x(t)$ that extremizes $S = \int L(x, \dot{x})\, dt$ with $L = \frac{1}{2}\dot{x}^2 - \Phi(x)$ (one spatial dimension, for clarity; the argument is identical in three dimensions with ∇ replacing d/dx), we use the Euler–Lagrange equation, which any calculus-of-variations or intermediate mechanics course derives from scratch:

$$\frac{d}{dt}\left(\frac{\partial L}{\partial \dot{x}}\right) - \frac{\partial L}{\partial x} = 0\,. \tag{9}$$

Plugging in. With $L = \frac{1}{2}\dot{x}^2 - \Phi(x)$:

$$\frac{\partial L}{\partial \dot{x}} = \dot{x} \quad \Longrightarrow \quad \frac{d}{dt}\left(\frac{\partial L}{\partial \dot{x}}\right) = \ddot{x}\,, \qquad \frac{\partial L}{\partial x} = -\frac{d\Phi}{dx}\,.$$

Substituting into Eq. (9):

$$\ddot{x} - \left(-\frac{d\Phi}{dx}\right) = 0 \quad \Longrightarrow \quad \ddot{x} = -\frac{d\Phi}{dx}\,.$$

In three dimensions, the identical steps (varying each Cartesian coordinate independently) give

$$\boxed{\ddot{\mathbf{x}} = -\nabla\Phi} \tag{10}$$

which is Newton's law of gravitation in its most general form: acceleration equals minus the gradient of the potential. For a point mass M at the origin, $\Phi(r) = -GM/r$, so $\nabla\Phi = (GM/r^2)\hat{\mathbf{r}}$ and

$$\ddot{\mathbf{x}} = -\frac{GM}{r^2}\hat{\mathbf{r}}\,,$$

pointing toward the mass, with the familiar inverse-square magnitude – exactly Newton's law, exactly as measured, exactly as taught since 1687. It required no assumption about the form of the force. The inverse-square law is not an input to this derivation; it falls out automatically once $\Phi = -GM/r$ is inserted, because that specific form of Φ is itself fixed independently by the requirement (not derived here, but standard) that Φ solve the Newtonian field equation $\nabla^2\Phi = 4\pi G\rho$ for a point source – this is simply Poisson's equation, the static, non-relativistic Gauss's law for gravity, with ρ the mass density; for a point mass $\rho = M\delta^3(\mathbf{r})$ and the solution is exactly $\Phi = -GM/r$. In other words, this musing derives *how* a free body responds to a given Φ; it borrows, rather than derives, *why* Φ takes the particular $1/r$ form it does.

7 The Other Half: Bending of Starlight

Everything up to this point leaned, quietly but constantly, on one approximation baked into Eq. (2): $v \ll c$. Light violates this maximally – $v = c$, exactly, always – so this is the place where the clock-only story must, honestly, meet its limit. It is also, happily, exactly where the piece this musing kept flagging as "discarded" finally has to reappear, and it reappears in a way fully within reach of the same tools.

7.1 Why Maximal Aging Cannot Govern Light

A photon's own proper time does not advance at all: $d\tau = 0$ along its worldline, identically, in any theory built on the Lorentzian interval – this is why a light clock, run at $v = c$, would read zero elapsed time no matter how far it travels. Maximal aging asked us to extremize τ over nearby paths; for light there is no τ to extremize. The correct replacement, standard in general relativity, is that light follows a **null geodesic**: a path along which the invariant spacetime interval itself vanishes at every point,

$$ds^2 = 0. \tag{11}$$

To use Eq. (11) we need the full interval ds^2, not merely the proper-time formula Eq. (5) we built it from – and that means we finally need the spatial part of the metric.

7.2 The One Ingredient This Musing Has to Import

Everything in Sections 2–6 used only the time-time piece, $g_{00} \approx -(1+2\Phi/c^2)$, and we obtained that piece honestly, from nothing but clocks in an accelerating rocket. The spatial metric components, g_{ij}, cannot be obtained the same way: a pure time-dilation argument, by its very construction, is blind to the geometry of space. Fixing g_{ij} requires solving the weak-field Einstein equations – a genuinely new input, not derived in this musing, imported here exactly as honestly as $\Phi = -GM/r$ was imported in Section 6. For a static, weak, spherically symmetric source, the standard result (see, e.g., [13], or [15] for a gentler, more physics-first derivation of the same weak-field metric) is

$$ds^2 = -\left(1 + \frac{2\Phi}{c^2}\right) c^2 dt^2 + \left(1 - \frac{2\Phi}{c^2}\right)\left(dx^2 + dy^2 + dz^2\right). \tag{12}$$

The time-time piece is exactly Eq. (4), reproduced intact – a welcome consistency check. The spatial piece is new: an *equal-magnitude, opposite-sign* correction to the geometry of space itself. This single extra term is the entire difference between this section and everything before it. (A reader who would like to know what $g_{\mu\nu}$ actually *is* – as a mathematical object, not just a symbol carried along since Eq. (12) – can turn to Appendix A, which builds it from scratch out of nothing more exotic than tilted coordinate axes.)

7.3 The Coordinate Speed of Light Near a Mass

Derivation: light slows down (in coordinate speed) near a mass. Consider a light ray moving radially, so $ds^2 = 0$ in Eq. (12) gives

$$\left(1 + \frac{2\Phi}{c^2}\right) c^2\, dt^2 = \left(1 - \frac{2\Phi}{c^2}\right) dr^2 \quad \Longrightarrow \quad \left(\frac{dr}{dt}\right)^2 = c^2\,\frac{1 + 2\Phi/c^2}{1 - 2\Phi/c^2} \approx c^2\left(1 + \frac{4\Phi}{c^2}\right)$$

to first order in Φ/c^2, so the coordinate speed of light is

$$v_{\text{light}}(r) \approx c\left(1 + \frac{2\Phi}{c^2}\right). \tag{13}$$

Since $\Phi < 0$ near a mass, $v_{\text{light}} < c$: light's *coordinate* speed (what a distant bookkeeper computes from dr/dt) drops near a mass, even though every local observer still measures exactly c for light passing by them. This coordinate slow-down is precisely the mechanism behind the Shapiro time delay, confirmed by radar-ranging experiments bouncing signals off Venus and off Viking landers on Mars.

7.4 Fermat's Principle: The Same Family of Idea, Once More

Equation (13) lets us treat the region around a mass exactly as an optical medium with a position-dependent refractive index $n(r) = c/v_{\text{light}}(r) \approx 1 - 2\Phi/c^2$. Since $\Phi < 0$ near a mass, $n(r) > 1$ there – exactly the familiar sense of "optically denser," the same sense in which glass or water has $n > 1$ and slows light down. This is not a loose analogy – geometric optics rays in a medium of varying n obey the same kind of extremal principle we have been using throughout: **Fermat's principle**, that light takes the path of least (more precisely, extremal) travel *time*, $\delta \int n\, ds/c = 0$. It is worth pausing on this: the light-bending calculation below is governed by exactly the same family of variational idea – extremize an elapsed time – as maximal aging was for the apple. Only the object being extremized (proper time for massive particles, optical path time for light) has changed.

The standard consequence of Fermat's principle for rays in a slowly-varying medium is the ray equation $\frac{d}{ds}\left(n\frac{d\mathbf{r}}{ds}\right) = \nabla n$, which produces a small transverse bend in an otherwise straight ray whenever n varies across the ray's path.

7.5 The Deflection Integral, and Why It Doubles

Derivation: deflection angle for a ray grazing a mass at impact parameter b. With $\Phi = -GM/r$, the full refractive index from Eq. (13) is

$$n(r) = 1 - \frac{2\Phi}{c^2} = 1 + \frac{2GM}{c^2 r}.$$

Treating the ray, to leading order, as an undeflected straight line at fixed impact parameter b (the standard weak-deflection approximation), parametrize points along it by $r = \sqrt{b^2 + x^2}$. The accumulated bending angle is the integral of the transverse gradient of n along the path:

$$\Delta\theta = \int_{-\infty}^{\infty} \frac{\partial n}{\partial b}\, dx, \qquad \frac{\partial n}{\partial b} = \frac{\partial n}{\partial r}\frac{\partial r}{\partial b} = -\frac{2GM}{c^2 r^2}\cdot\frac{b}{r} = -\frac{2GMb}{c^2 r^3}.$$

Using $\int_{-\infty}^{\infty} dx/(b^2 + x^2)^{3/2} = 2/b^2$:

$$\Delta\theta_{\text{full}} = \frac{2GMb}{c^2}\cdot\frac{2}{b^2} = \frac{4GM}{c^2 b}. \tag{14}$$

Now repeat the identical calculation, but with only the time-time curvature restored and space left flat – exactly the "time-only" story this whole musing has otherwise told. Setting $g_{ij} = \delta_{ij}$ in Eq. (12) and redoing the null-geodesic derivation above gives $v_{\text{light}}(r) \approx c(1 + \Phi/c^2)$ – half the coefficient of Eq. (13) – hence $n_{\text{time-only}}(r) = 1 + GM/(c^2 r)$, exactly half the full $n(r)$. The bending-angle integral is linear in $n - 1$, so

$$\Delta\theta_{\text{time-only}} = \frac{2GM}{c^2 b} = \frac{1}{2}\,\Delta\theta_{\text{full}}\,. \tag{15}$$

Equation (15) is the payoff. *Time curvature alone – the entire content of Sections 2 through 6 – accounts for exactly half of the true bending of starlight; spatial curvature, deliberately absent from this musing until now, supplies the other half.* This is not a coincidence of algebra; it is the general lesson of the weak-field metric, Eq. (12): for slow-moving matter the g_{ij} term is invisible (it only multiplies v^2/c^2, already second order), but for light, moving at $v = c$, the two curvatures contribute on exactly equal footing.

History got here in exactly this order. Einstein's 1911 paper [2], using only the equivalence principle – the time-only story of Sections 2–6 of this musing – predicted $\Delta\theta = 2GM/(c^2 b)$: Eq. (15). This number coincides with an even older, purely Newtonian calculation treating light as a fast corpuscle, first carried out by von Soldner in 1804 [3]. Only in 1915, with the full field equations in hand [4], did Einstein obtain the doubled result, Eq. (14). Eddington's 1919 eclipse expedition [11] measured the deflection of stars near the darkened sun and confirmed the doubled value, not the halved one – the single most famous observational vindication of general relativity over any time-only or purely Newtonian account, and a direct, historical confirmation that spatial curvature is not optional once $v \to c$.

Numerical check: the sun. For the sun, $GM_\odot/c^2 \approx 1.475$ km (the same combination that sets the Schwarzschild radius, $2GM_\odot/c^2 \approx 2.95$ km). Grazing starlight has impact parameter $b \approx R_\odot \approx 6.96 \times 10^5$ km. Then

$$\Delta\theta_{\text{full}} = \frac{4 \times 1.475}{6.96 \times 10^5}\ \text{rad} \approx 8.48 \times 10^{-6}\ \text{rad} \approx 1.75\ \text{arcsec}\,,$$

using 1 rad = 206,265 arcsec – the exact number Eddington's team measured in 1919, and exactly twice the 0.87 arcsec that Einstein's 1911 time-only estimate, or Soldner's 1804 corpuscular one, would have predicted.

8 Where Time Stops: The Black Hole Horizon

Every section so far has trusted the same clock. It ticked at a slightly different rate on an airplane. It ticked faster on a GPS satellite and slower on the ground. It ticked more slowly, by exactly the amount needed, deep in a well of gravitational potential – and it was that one small effect, taken seriously and pushed through the calculus of variations, that produced Newton's law of gravity out of nothing more exotic than a rocket-and-photon thought experiment. This section asks the last question that clock can answer: what happens if you keep lowering it? Not to the surface of the Earth, or the surface of the sun, but deeper – into a well so steep that the gentle, linear approximation this musing has leaned on since Section 2 simply stops being gentle. The clock does not do anything qualitatively new. It just keeps doing the same thing, more and more, until there is nowhere left for "more" to go.

Section 7 pushed past one approximation ($v \ll c$) while keeping another firmly in place: everywhere in this musing, Φ/c^2 has been treated as small, and every formula has been truncated after the first power of it. That is an excellent approximation for the sun ($GM_\odot/(Rc^2) \sim 10^{-6}$) but it is, by construction, blind to what happens when a potential well becomes so deep that Φ/c^2 is no longer small at all. This section drops the linearization – the one remaining approximation – and asks what the exact, un-truncated time-dilation formula predicts.

8.1 The Exact Formula, No Taylor Expansion

For a static, spherically symmetric mass M, the exact solution of the full Einstein field equations – not the linearized version used everywhere above – is the Schwarzschild metric [12]. Written in the standard radial coordinate r (a different, and for this purpose more convenient, coordinate choice than the isotropic-style coordinates behind Eq. (12) in Section 7 – both are standard weak-field coordinate choices, and both describe the same physical spacetime, agreeing once expanded to first order in Φ/c^2), its time-time and radial-radial components are

$$g_{00} = -\left(1 - \frac{2GM}{rc^2}\right), \qquad g_{rr} = \left(1 - \frac{2GM}{rc^2}\right)^{-1}. \tag{16}$$

Compare g_{00} here to Eq. (4): with $\Phi = -GM/r$, our weak-field result was $g_{00} \approx -(1 + 2\Phi/c^2) = -(1 - 2GM/rc^2)$ – exactly the first two terms of Eq. (16)'s g_{00}, since $(1 - 2GM/rc^2)$ is already linear in GM/rc^2 with no higher corrections to drop. The clock-derived formula from Section 2.2 was not an approximation to the wrong thing; it was the exact leading behavior of the right thing, valid as far as its own derivation (a weak-field rocket argument) could honestly reach. What changes now is that we take Eq. (16) as exact and stop assuming GM/rc^2 is small.

A natural length scale falls out immediately. Equation (16) is singular where $1 - 2GM/rc^2 = 0$, i.e., at

$$r_s \equiv \frac{2GM}{c^2}, \tag{17}$$

the **Schwarzschild radius**. Nothing about this musing's earlier sections anticipated a special radius – it emerges here purely because we stopped discarding higher-order terms in Φ/c^2.

8.2 A Clock's Tick Rate, Exactly

Derivation: exact gravitational time dilation. A clock sitting still at radius r (constant r, θ, ϕ) traces out a worldline with $ds^2 = g_{00}c^2dt^2 = -\left(1 - \frac{2GM}{rc^2}\right)c^2dt^2$. Since $ds^2 \equiv -c^2d\tau^2$ by definition of proper time, this gives immediately

$$\frac{d\tau}{dt} = \sqrt{1 - \frac{r_s}{r}}. \tag{18}$$

No Taylor expansion has been used anywhere in this derivation – Eq. (18) holds for any $r > r_s$, however deep the well.

Equation (18) is the exact version of the same statement Section 2.2 made approximately: a clock deeper in a potential well runs slow. Lower it a little, and it loses a little time. Lower it further, and it loses more. There is nothing dramatic in this yet – it is the identical arithmetic behind the GPS correction in Section 2.2, just no longer stopped after its first term. But now watch what happens as r keeps shrinking toward r_s: the clock does not level off, or approach some small floor. It keeps slowing, and slowing, and slowing, with no bottom in sight, until, at $r \to r_s^+$, $d\tau/dt \to 0$. **A clock hovering just above the Schwarzschild radius ticks, from the point of view of a distant bookkeeper, arbitrarily slowly – and at $r = r_s$ itself, it appears to stop entirely.** This is not a new postulate, and nothing new has been assumed about clocks between here and page one. It is the same idea, run to its logical end.

8.3 The Coordinate Speed of Light, Exactly

Repeating the null-geodesic calculation of Section 7.3, but now with the exact Eq. (16) rather than the linearized Eq. (12), for a radially infalling or outgoing light ray ($d\theta = d\phi = 0$): setting $ds^2 = 0$ gives

$$\left(1 - \frac{r_s}{r}\right) c^2 dt^2 = \frac{dr^2}{1 - r_s/r} \quad \Longrightarrow \quad \left(\frac{dr}{dt}\right)^2 = c^2 \left(1 - \frac{r_s}{r}\right)^2 ,$$

so the exact coordinate speed of light is

$$v_{\text{light}}(r) = c\left(1 - \frac{r_s}{r}\right) , \tag{19}$$

which vanishes exactly at $r = r_s$: to a distant bookkeeper using coordinate time t, light emitted radially outward from the horizon itself never gets anywhere. This is the coordinate-speed statement whose linearized cousin, Eq. (13), only ever showed light *slowing*; the exact formula shows it *stopping*.

A caveat, stated honestly. The vanishing in Eq. (19) and the divergence of $1/g_{rr}$ at $r = r_s$ are properties of *this particular coordinate system* (Schwarzschild coordinates), not of the physical spacetime itself: an observer who actually falls through $r = r_s$ experiences nothing special there – finite tidal forces, finite proper time to cross, no local sign of a boundary. Switching to horizon-regular coordinates (Eddington–Finkelstein, or Kruskal–Szekeres; not pursued here, see [13]) removes the coordinate singularity entirely while keeping r_s itself as a genuine, coordinate-independent one-way membrane: light and matter cross inward but never back out. The freezing in Eqs. (18)–(19) is real for the distant bookkeeper – it is exactly what a distant observer actually measures – but it is not evidence that anything catastrophic happens *locally* at the horizon itself.

Numerical check: three Schwarzschild radii. Using $r_s = 2GM/c^2$ and the same $GM_\odot/c^2 \approx 1.475$ km from Section 7:

- **The sun**: $r_s = 2 \times 1.475$ km ≈ 2.95 km – about the size of a small asteroid, buried deep inside the sun's actual radius of 6.96×10^5 km, which is exactly why the sun is not a black hole: nothing compresses its mass inside r_s.
- **The Earth**: $M_\oplus/M_\odot \approx 3.0 \times 10^{-6}$, so $r_s \approx 2.95$ km $\times\, 3.0 \times 10^{-6} \approx 8.9$ mm – Earth's entire mass would need to fit inside a radius smaller than a grape.
- **Sagittarius A***, the supermassive black hole at the center of the Milky Way, $M \approx 4.15 \times 10^6\, M_\odot$: $r_s \approx 2.95$ km $\times\, 4.15 \times 10^6 \approx 1.22 \times 10^7$ km ≈ 0.082 AU – smaller than Mercury's orbit, yet this is the actual, physical, one-way boundary of an object whose

shadow the Event Horizon Telescope imaged directly in 2022.

9 What Was Not Needed for the Apple, and Why That Mattered

It is worth totalling up, at the end, exactly what machinery this derivation used and what it deliberately avoided.

The complete ingredient list.

- Special relativity's time-dilation formula for a moving clock (1905), Eq. (1).
- The equivalence principle, applied to an accelerating rocket (1907), giving gravitational time dilation, Eq. (4).
- A single Taylor expansion to combine the two effects, Eq. (5).
- One physical postulate: free particles maximize elapsed proper time ("maximal aging"), Section 4.
- Ordinary calculus of variations – the Euler–Lagrange equation – applied to the resulting integral, Section 6.

Not used, for the apple: the metric tensor $g_{\mu\nu}$, the Christoffel symbols $\Gamma^{\mu}_{\alpha\beta}$, the Riemann curvature tensor, the Einstein field equations, or any concept of curved space. Every one of those objects is real, essential to relativity, and completely necessary the moment speeds approach c or fields become strong. Section 7 showed exactly this happening for light: restoring the spatial metric g_{ij} – deliberately withheld through Sections 2–6 – exactly doubled the deflection of starlight, from the time-only value Einstein predicted in 1911 to the correct value confirmed at the 1919 eclipse. Section 8 showed the second boundary of ordinary experience: dropping the linearization in Φ/c^2 entirely, rather than the speed approximation, revealed a radius at which a clock's tick rate falls to zero – the black hole horizon, obtained from the very same time-dilation formula as Section 2.2, simply carried past the point where truncating it after one term stops being harmless. But for the apple, the falling ball, the orbiting satellite, and every projectile Newton ever considered, none of that machinery was required. Gravity, at the speeds and scales of ordinary experience, is entirely a statement about clocks – and light and black holes are precisely the two phenomena that show you, in two different directions, where "ordinary experience" stops.

The rubber sheet was never wrong because curvature is a bad concept – embedding diagrams of curved spatial slices are a legitimate and standard tool in GR (they are exactly how one visualizes, say, the spatial geometry around a black hole, or why light rays bend near the sun once both time and space curvature are included, as Section 7 just showed explicitly). The rubber sheet is wrong specifically as an explanation of *why an apple falls*: it is wrong because it curves the wrong thing (space, not time) for that particular phenomenon, for the wrong reason (it visibly needs gravity, via the ball's weight, to produce the very curvature meant to explain gravity), and because it invites the viewer to conflate a picture of the spatial slice with the actual 4-dimensional geometry that is doing the physical work. The wristwatch never makes any of these three mistakes for the apple – and even for light, where the wristwatch runs out of road, it hands off cleanly to Fermat's stopwatch, extremizing optical path time by exactly the same variational logic; and even at a horizon, where a distant clock appears to stop altogether, the story is still the same wristwatch, pushed past the point of gentle approximation rather than replaced by a new idea. Everything in this musing – apple, starlight, and horizon alike – is the arithmetic of a system trying to extremize the time it takes – proper time for matter, optical time for light – and Newton's law, like the 1919 eclipse result and the freezing of a clock at r_s, is simply what that arithmetic works out to. Appendix A supplies the one piece of scaffolding assumed throughout but never explained: what the metric tensor $g_{\mu\nu}$, quietly doing all this work since Eq. (12), actually *is*.

Gravity did not begin, in this telling, as a force. It began as time refusing to pass equally everywhere.

A A Primer on Metrics: Tilted Axes, and Why Components Come in Two Kinds

Sections 7 and 8 leaned on objects called g_{00}, g_{ij}, $g_{\mu\nu}$ without ever explaining what a metric tensor actually *is*. This appendix fills that gap, starting from the most familiar picture possible – ordinary coordinate axes – and asking what changes if those axes are allowed to be non-orthogonal and of unequal length.

A.1 The Familiar Case: Orthonormal Axes

In ordinary Cartesian coordinates, the axes $\mathbf{e}_x, \mathbf{e}_y$ are unit length ($|\mathbf{e}_x| = |\mathbf{e}_y| = 1$) and mutually perpendicular ($\mathbf{e}_x \cdot \mathbf{e}_y = 0$). A displacement $d^2 = dx^2 + dy^2$ has this simple Pythagorean form precisely *because* of these two properties. There is also, in this familiar case, no distinction worth making between "how far you went along each axis" and "the shadow (projection) the displacement casts on each axis" – they are the same number, for each axis, because the axes are orthonormal.

A.2 Tilting the Axes

The question this section answers. Suppose the two axes $\mathbf{e}_1, \mathbf{e}_2$ are still unit length, but the angle between them is some $\theta \neq 90°$ rather than a right angle. A given displacement $\mathbf{d}$ can now be described in two genuinely different ways. What are they, and how are they related?

Way 1 – contravariant components (the parallelogram rule). Write $\mathbf{d} = A\,\mathbf{e}_1 + B\,\mathbf{e}_2$: go A steps along $\mathbf{e}_1$, then B steps along $\mathbf{e}_2$, and complete the parallelogram. (A, B) are the **contravariant components** of $\mathbf{d}$, conventionally written with an upper index, $d^1 = A$, $d^2 = B$.

Way 2 – covariant components (projections). Alternatively, describe $\mathbf{d}$ by its projection (shadow) onto each axis: $d_1 \equiv \mathbf{d} \cdot \mathbf{e}_1$, $d_2 \equiv \mathbf{d} \cdot \mathbf{e}_2$. These are the **covariant components**, written with a lower index.

The mnemonic to keep. Upper index = contravariant = parallelogram-rule components ("how far along each axis"). Lower index = covariant = projection components ("the shadow on each axis"). The two agree, index-for-index, only when the axes are orthonormal – which is exactly why ordinary Cartesian calculus never needs to distinguish them.

Deriving the relationship. With $\mathbf{d} = A\mathbf{e}_1 + B\mathbf{e}_2$,

$$d_1 = \mathbf{d}\cdot\mathbf{e}_1 = A(\mathbf{e}_1\cdot\mathbf{e}_1) + B(\mathbf{e}_2\cdot\mathbf{e}_1)\,, \qquad d_2 = \mathbf{d}\cdot\mathbf{e}_2 = A(\mathbf{e}_1\cdot\mathbf{e}_2) + B(\mathbf{e}_2\cdot\mathbf{e}_2)\,.$$

Define the **metric tensor** as simply the table of all pairwise dot products of the basis vectors, $g_{ij} \equiv \mathbf{e}_i \cdot \mathbf{e}_j$. For our unit-length axes at angle θ: $g_{11} = g_{22} = 1$ (unit length) and $g_{12} = g_{21} = \cos\theta$ (the one new number, zero exactly when the axes are orthogonal). Then the two boxed equations above are simply

$$d_i = g_{ij}\, d^j \qquad \text{(sum over } j)\,, \tag{20}$$

the operation conventionally called **lowering an index**: the metric tensor is nothing but the dictionary between the parallelogram-rule description and the projection description of the same vector.

Worked numerical example. Let $\theta = 60^\circ$, so $\cos\theta = 0.5$, and take a displacement with contravariant components $(d^1, d^2) = (2, 1)$ – two steps along $\mathbf{e}_1$, one step along $\mathbf{e}_2$. Then, from Eq. (20):

$$d_1 = g_{11}d^1 + g_{12}d^2 = (1)(2) + (0.5)(1) = 2.5\,, \qquad d_2 = g_{21}d^1 + g_{22}d^2 = (0.5)(2) + (1)(1) = 2\,.$$

Notice $d_1 \neq d^1$ and $d_2 \neq d^2$: covariant and contravariant components genuinely differ once the axes are non-orthogonal, even though (d^1, d^2) and (d_1, d_2) describe exactly the same physical displacement $\mathbf{d}$. The actual physical (squared) length of $\mathbf{d}$, which must not depend on which description we use, is computed as $ds^2 = g_{ij}d^i d^j = (1)(2)^2 + 2(0.5)(2)(1) + (1)(1)^2 = 4+2+1 = 7$, i.e. $|\mathbf{d}| = \sqrt{7} \approx 2.646$ – the same answer either bookkeeping method must give, and does.

A.3 Back to Spacetime

The metric $g_{\mu\nu}$ used throughout this musing (Eq. (5) implicitly, Eq. (12) and Eq. (16) explicitly) is exactly this same object – a table of basis-vector dot products, $g_{\mu\nu} = \mathbf{e}_\mu \cdot \mathbf{e}_\nu$ – just built from a four-dimensional coordinate grid (t, x, y, z) instead of a tilted two-dimensional one, and using the Lorentzian rule $\mathbf{e}_0 \cdot \mathbf{e}_0 = -1$ instead of $+1$ for the time axis. This sign is purely a convention, not a physical fact – this musing follows the "mostly plus" signature $(-, +, +, +)$ used by, e.g., [14, 15]; a great deal of the GR literature, including [13], instead uses "mostly minus" $(+, -, -, -)$, with every formula in this musing correspondingly flipped in overall sign but physically identical. In flat, empty spacetime with no mass anywhere, this table is the constant diagonal matrix $\mathrm{diag}(-1, 1, 1, 1)$ (the Minkowski metric) – the four-dimensional analogue of orthonormal Cartesian axes. What Eq. (12) says is that near a mass, the basis vectors' *lengths* change from point to point (g_{00} and g_{ij} both depart from ± 1), even though this particular weak-field metric happens to stay perfectly *diagonal* everywhere – the axes stretch and shrink, but (in these coordinates) never tilt into one another. This is the precise, technical sense in which spacetime is "curved" near a mass: not that the coordinate grid is tilted (it isn't, here), but that the ruler you use to measure distances along that grid changes length from one point to the next, in a way that cannot be undone by any single, global relabeling of the axes. This musing never needed to raise or lower an index explicitly, because every physical quantity used ($d\tau$, ds^2) was built directly as $g_{\mu\nu}\, dx^\mu dx^\nu$ rather than manipulated component-by-component – but the dictionary of Eq. (20) is exactly what is quietly doing the work every time g_{00} or g_{ij} appears. A reader who wants the fully developed machinery – the Christoffel symbols $\Gamma^\mu_{\alpha\beta}$ that appear when the axes' *tilt* itself changes from point to point, and the covariant derivative built from them – can find a careful undergraduate-level treatment in [14], and the complete graduate-level treatment in [13].

References

[1] A. Einstein, "Über das Relativitätsprinzip und die aus demselben gezogenen Folgerungen," *Jahrbuch der Radioaktivität und Elektronik* **4**, 411–462 (1907). The paper containing the first version of the equivalence principle and gravitational time dilation/redshift argument used in Section 2.2.

[2] A. Einstein, "Über den Einfluss der Schwerkraft auf die Ausbreitung des Lichtes," *Annalen der Physik* **35**, 898–908 (1911). The time-only estimate of light bending, using the equivalence principle alone, that Section 7 reproduces exactly as Eq. (15).

[3] J. G. von Soldner, "Über die Ablenkung eines Lichtstrals von seiner geradlinigen Bewegung," *Berliner Astronomisches Jahrbuch*, 161–172 (1804). The purely Newtonian, pre-relativistic corpuscular-light calculation that coincidentally matches Einstein's 1911 time-only estimate.

[4] A. Einstein, "Die Grundlage der allgemeinen Relativitätstheorie," *Annalen der Physik* **49**, 769–822 (1916). The full field-equations calculation giving the doubled deflection, Eq. (14).

[5] E. F. Taylor and J. A. Wheeler, *Exploring Black Holes: Introduction to General Relativity*, Addison Wesley Longman (2000). The source of the "principle of maximal aging" framing and the parable-of-two-travelers style arithmetic comparison used in Section 4.

[6] J. C. Hafele and R. E. Keating, "Around-the-World Atomic Clocks: Predicted Relativistic Time Gains," *Science* **177**, 166–168 (1972).

[7] J. C. Hafele and R. E. Keating, "Around-the-World Atomic Clocks: Observed Relativistic Time Gains," *Science* **177**, 168–170 (1972).

[8] N. Ashby, "Relativity in the Global Positioning System," *Living Reviews in Relativity* **6**, 1 (2003). Source for the +38 μs/day net GPS clock drift figure quoted in Section 2.2.

[9] P.-L. Moreau de Maupertuis, "Accord de différentes lois de la nature qui avaient jusqu'ici paru incompatibles," *Mémoires de l'Académie Royale des Sciences* (1744).

[10] W. Yourgrau and S. Mandelstam, *Variational Principles in Dynamics and Quantum Theory*, 3rd ed., Dover (1968). A readable history of the least-action principle from Maupertuis through Euler and Lagrange.

[11] F. W. Dyson, A. S. Eddington, and C. Davidson, "A Determination of the Deflection of Light by the Sun's Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919," *Philosophical Transactions of the Royal Society A* **220**, 291–333 (1920). The eclipse measurement confirming the full (space-plus-time-curvature) deflection is twice the naive Newtonian/time-only estimate.

[12] K. Schwarzschild, "Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie," *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, 189–196 (1916). The original exact solution used in Section 8, Eq. (16).

[13] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*, W. H. Freeman (1973). Standard reference for the weak-field metric, the PPN factor-of-2 light-bending calculation, the geodesic equation cited in the Section 4 conceptbox, and horizon-regular (Eddington–Finkelstein/Kruskal) coordinates mentioned in Section 8.

[14] B. F. Schutz, *A First Course in General Relativity*, 2nd ed., Cambridge University Press (2009). An accessible undergraduate-level development of the tensor calculus (covariant derivatives, Christoffel symbols) sketched only informally in Appendix A.

[15] J. B. Hartle, *Gravity: An Introduction to Einstein's General Relativity*, Addison Wesley (2003). "Physics first" pedagogy discussed in the companion conversation to this musing; introduces the Schwarzschild metric and weak-field results well ahead of the field equations.