Mechanics — Narrative

Executive recap (our picture, no fluff)

Mass is missing space: a closed, light-speed circulation that keeps a fixed slice of space hidden. Bring another spinning loop near it and the clean route leans slightly toward that hidden-space profile and takes a hair longer to pass. That's gravity: route in a shaped background set by another loop's circulation. No "pulls." Same two rules as always: no sideways pile-up across boundaries; no tearing of closure.

Scope for this pillar (what we do and what we don't)

Mechanics carries the full gravity story in VMS — from gentle bends to horizon-scale effects. We keep the same ontology (routes/closure, display-area $A_d = \pi r^2$), the same bookkeeping ($S[path] = \int A_d\, ds$), and the same single scale $S_0 = \hbar$ (SI-locked).

• In-scope: weak-field (tiny bend, small delay, clock drift) and strong-field (deep-well lensing, photon-sphere/horizon behavior, periapsis precession in steep wells, frame-dragging around rotators, wave transport across steep gradients).

• Checks: recover the standard weak-field numbers and the tested strong-field signatures using the same settings; no new tunable scales.

• Out-of-scope here: cosmology/expansion; quantum corrections beyond S_0 (Particle pillar); heavy numerics beyond the kernels shown in the Walk-Through.

What the Math Walk-Through Will Actually Do

We start with the same bookkeeping you've seen everywhere in this pillar. A Void advancing at light speed hides a circular patch of space as it moves. That patch is the display-area $A_d = \pi r^2$. Along any route, we simply add up how much area gets hidden at each step. That running total is the route 'cost' $S[path] = \int A_d\, ds$. That running total is the route 'cost' **(display-area action)** $S[path] = \int A_d\, ds$. At first we only compare ratios of S—cleaner vs. costlier routes—no new dials, no imported forces.

A stationary mass is a loop that keeps a chunk of space hidden—missing space. From a distance, the detailed swirl blurs into a gentle background that says "steps here cost a tiny bit more." The Walk-Through names that background n(x). Think of **n(x)** as a smooth *cost map*: higher n means 'this step hides a bit more display-area than average.' It's not a new substance. It's just a smooth way to keep track of the missing-space profile in our own language.

Once you have that background, tiny sideways changes in it do three simple things to passing routes:

• They tip the route slightly toward the mass (a small bending angle).

• They make near passes take a hair longer than far passes (a small positive delay).

• They make parked loops deep in the profile tick a little slower than ones far away (a small clock drift).

The Walk-Through shows each of these with one compact expression—little path integrals you can actually plug into. Same knobs as always: one scale $S_0 = \hbar$ (SI-locked), c and G as acceptance locks, everything else is geometry. No extra parameters appear halfway through the derivation.

How the classical numbers show up (why we're compatible)

After the story above, we take one smoothing step: blur the rotating loop into a background and keep the leading terms. In that weak-field limit, the three outputs—bend, delay, and clock drift—land on the familiar textbook values (the usual deflection, Shapiro-style travel-time shift, and gravitational redshift/time-dilation). We don't import someone else's field equations as axioms; we just show that our route picture reduces to the same tested numbers when you look from far away. We set the matching coefficients once (those coefficients are the standard weak-field factors used across deflection, delay, and clock-shift) and then reuse them across every case.

Strong-Field Narrative (what changes and why)

All the pieces are the same — only the gradients get big. A stationary mass is still a closed circulation that keeps space hidden (missing space). From far away the swirl blurs to a smooth background cost n(x); near compact masses that background gets steep. The same two rules hold: no sideways pile-up across boundaries; no tearing of closure. Those rules plus a steep n(x) profile give you:

• Routes that wind hard: light and matter bend strongly; near the photon-sphere the clean options wrap before they escape.

• Passes that get expensive: near-mass paths pay a large time premium (deep-well delay); parked clocks drift heavily (gravitational time-dilation).

• Background that twists if the mass rotates: the route frame is dragged — you see it as a small azimuthal shove (frame-dragging/Lense–Thirring in smooth limit).

Nothing new was imported. We're still counting hidden display-area along routes. Strong-field just means the cost landscape is steeper, so the same choice rule (cleaner route wins) produces stronger bending, bigger delays, and deeper clock shifts.

Where the Walk-Through will take it (preview)

• Write the strong-field kernels as the same path integrals you already saw in weak-field, now evaluated on steep profiles n(x).

Plain-language checklist (use this as you read the math)

• A_d and S are the only moving parts at the start; n(x) is just a smooth summary of the missing-space profile.

• Three consequences, one cause: tiny sideways gradients in that profile.

• Outputs are small, positive-defined effects you can integrate along a path (no new dials appear).

• The weak-field bridge exists to prove we line up with the known numbers—once—then we move on.

• "Show how the dictionary to familiar GR results appears in the smooth limit (deflection near compact masses, deep-well timing, frame-dragging terms)."

• "Keep the locks: $S_0 = \hbar$ once; c and G as acceptance locks; no mid-derivation knobs."

Bench protocols (what you can actually run)

• Cavity frequency pull near mass: park a high-Q cavity near a stationary loop (mass) and read the frequency shift; compare to the delay kernel.

• Ring-laser/fiber-gyro bias/null: place mass symmetrically; beyond Sagnac, expect a null. Any bias sets a bound.

• Bench-top delay: near-mass pass vs far-pass timing difference (Shapiro analog at lab scale).

• Polarization near mass: expect no rotation/birefringence in this bridge limit; any detection is a bound.

Falsifiers (clean failure modes)

• Chromatic gravity in the bridge limit (wavelength-dependent bend/delay) would be a fail or an extensions lead.

• Polarization rotation near mass in symmetric placement would be a fail or a bound.

• Non-reciprocity beyond Sagnac in symmetric rings would be a fail or a bound.

Language & locks

• Projected area only: $A\_d = \pi r^2$. No 4π shells. No double-counting.

• Single scale: $S_0 = \hbar$ (fixed via SI's exact h); it's the converter from action to phase when we need phase.

• No extra dials: we don't retune per experiment; predictions live on the same yardstick.

• No flux talk: we stay in routes/closure terms the whole time.

Reader map

• Walk-Through — the compact derivations for bend, delay, and clock drift; smooth-limit bridge included.

• Math Appendix — the readable equations and validity limits (weak-field, paraxial, symmetry notes).

• Calibration & Verification — procedures, pass/fail tables, ratio-first when tighter than absolute calibration.

• Student Workbook — pedagogy and practice problems (separate track).

• VMS Laws — concise value cards (one page per law) for quick use.

Bottom line

We keep the picture simple and the promises tight: one ontology, one scale ($S_0 = \hbar$), three gravity outcomes. The Walk-Through turns that into short formulas you can actually use; the Appendix and tests keep us honest.