

Loop Walking Cascadia: a Triangle Lattice, its Rogues and Witnesses, and One Opening of the Vault

Matthew “Bear” Schulz

Independent Researcher · Belmont, CA · ORCID 0009-0000-0956-5743 ·

mattschulz97@protonmail.com · License CC BY 4.0 · Version 1.0 · 2026-08-24

“Loud isn’t equal to effective. That lesson the stars and quiet quasars taught me.”

Author’s note

The shape of this paper is a saddle — a two-part story. We will share our original plan, our findings and results, and our post-hoc adventures. And the most interesting result of the first half was a new, final cliffhanger question from section 10.4: what does winding mean? The question we could not close. Now, that is way too tempting a cliffhanger for me, and I finish a paper not expecting to ever revisit it. So Cascadia, Part Deux builds the saddle as we cowboy up for the last ride of Cascadia’s Holonomy Cowboys.

We refer to our first efforts as the conventional geodesy direction of playing skiball. In the later half, starting at the Interlude, we change the game to pachinko. Why? Inspiration came in the form of an imagined glass full of ice as we fill the glass with fizzy water and play pachinko inside the earth. And that was the twist the Cascadia data needed — both the visualization of the pachinko pins mechanism and its results.

§0 — Hello, World of Geodesy

Data needs to speak, and it speaks through visualization and metaphor. Analogy is not decoration here; it is the method’s front door. So let me tell you how a poker table led to a subduction zone.

My introduction to holonomy was born chasing adversarial liars around a poker table. Each player — each sensor — was a quaternion-driven model of each hand, and the winning hand’s journey, plotted first quaternion to last, closed into a loop. The cards were the sovereign substrate. The other players were the Byzantine sensor network, tilting perception of that substrate to reprice information. At about the same time, invariant-based, radiation-hardened memory systems fell out nearly fully formed from my Artificial Life work. Those two projects introduced me to holonomy.

When I discovered holonomy I was awestruck by its quiet power. Many disciplines use it, but it isn’t taught as a methodology. Practical holonomy is not a thing. It is the Method of the Loop Walker — differential geometry and curvature are its native language — and my first goals were to make holonomy accessible, teach folks across domains how it manifests, where it shines, and where it goes blind with magnitude. Why is it so little known? It’s a physics thing, and the example is always the same arrow-walk on a globe demonstrating curvature. Go to the holonomy page on Wikipedia and you’ll find that arrow walk and a whole universe of theoretical manifolds beneath it. In several thousand hours of this work I have used one 3-manifold, once. Everything else is built on that maybe-too-simple arrow walk at the top. If you don’t get it right away, welcome to the club — I struggled.

Drift is usually what you hunt and eliminate; invariants are what you keep. Holonomy is the instrument that tells them apart: walk a loop, and whatever failed to come home unchanged has named itself.

My earlier road into the earth ran through Euler: model the crust as imperfect spherical plates floating on the mantle, each junction and seam holding translation, motion, and curvature — model each seam as the junction of two or more Euler qualities. The junctions and their angles store the information; the seam itself loads that strain gauge. My initial success with holonomy and geodesy caused me to sprint overconfidently into the graveyard that has welcomed geodesy people since the dawn of data tracking. During the Topological Tripwire paper (concept DOI 10.5281/zenodo.21287950) I had found some interesting squiggles that could be a precursor shadow. Now, I think it’s human nature to want to find something cool. So I had my shiny portable holonomy instrument, and wondered if the first lock it might open was a different way to examine precursors. But instead of patience and care, I dove in headfirst, wickedly overconfident. I was ready to TELL that data how, when, and where it was going to speak to me on MY geometric terms. I brought my cleverest ideas of geometry to bear on the deep Earth. And I got clobbered. But failure is the best teacher. Always fail forward.

I wanted to force my geometry: what would be the *best* n-gon loop to capture the local curvature of strain and flexion between plates and in junctions? That sinking feeling came when I realized I was forcing my geometry on the data instead of listening to it. I hadn’t validated any geometric loop model. I didn’t know if loop size or location mattered.

Do you walk the plates' junction, or across the seams? Does the answer change if the n-gon misses the epicenter? That's when the little rules started to bend in productive ways — the rules were assumptions and blind spots I'd brought with me.

What happens when I assume a pentagon? The curvature tells me if I guessed right: the +1 pentagon gravity well that reports as a -1 heptagon is truly a heptagon in that moment, regardless of the n-gon I drew. And the special sauce is the hexagon: hexagons tile flat, so a zero-curvature patch reads as a hexagon, and every zero-reading seam is a free calibration point.

So I knew I was ill-prepared to continue the geodesy work, because I had just proven the need for a significant methods validation exercise before returning to geodesy. I went and chased DNA, knots, metallurgy, and built the Log and Anti-Log Towers of the Integer Forest — a pure mathematical energy-network map validated through the first hundred billion integers, its octave laws and pure structure standing in service of multiplicative number theory. That detour calibrated the tool on pure mathematical structure and gave us a coordinate system. So when I came back, I came back humbler. Why not take a break hunting white whales and do something fundamentally productive? Let's REALLY listen and visualize what the DATA wanted to share. One stops asking what loop to draw and asks instead: what is the optimum n-gon pattern the DATA recommends? That's when we modeled the GPS stations as a fully interconnected network node map — every station checked and connected against the others. Why fully interconnected? Because the universe is a wonderfully interconnected space, and holonomy gives us access to its structural wiring as fiber and sphere bundles, pseudo or real. All we do is hack that substrate by splicing into the network and walking the loop.

So, I needed to select an active, data-rich environment for the triangle lattice. Japan is the perfect, mature candidate for the next campaign. However, I needed to prove it out back in the Cascadia box. For a lattice of logarithms overlaid with triangles, it would push boundaries and be a good teacher. Cascadia would stubbornly refuse ANY attempt at my enforced geometric ideas by design. While Japan has a network that is perfect to combine lattices and differential-geometry loop walks, I was comfortable playing with Cascadia's data fields.

Here is where that question finally dissolved, and I love this: the old question assumed the signal lives in one well-chosen loop, and placement was the design problem — the original Parkfield pentagon and the load-share ring were that thinking, and they were right for their moment. What the triangle lattice did is dissolve the placement problem. Because the lattice is oriented and discrete Stokes holds, machine-verified, every possible n-gon is just a sum of triangles — the 9-gon, the 200-gon, a nested ring, the boundary of a sub-block, the fan around one suspect station. You don't place loops in advance anymore; you compute all 441 triangles once, and any loop becomes a query you ask afterward. The perfect n-gon isn't found — it's *selected*, as the boundary of whatever triangle set your question picks out. And the thing that replaced placement as the real object is the interconnectivity itself and the common-quality states: the kinematic charts turned out to be first-class physics, not preprocessing; the seams between them are where the loud signals live; the block-versus-block matrix is the coarse lattice over the fine one. So the old question — where do we put the loop? — dissolved into its answer: measure the whole fabric once, and afterward any loop you care to ask about is just a query against it.

§1 — The Instrument: from Loop Walks to a Triangle Lattice

The pipeline is four stages, each one earning the next. Stage zero grades every station in the public Nevada Geodetic Laboratory archive for the Cascadia box and issues quality cards; stations die only for stated reasons, and nothing dies without one — the cemetery rule. Stage one asks who dances with whom: each station's detrended series is its dance, and stations that dance together group together regardless of distance. Run that clustering twice — once kinematic (who moves alike), once geographic (who sits near whom) — and the disagreements between the two maps are themselves findings: coherent crustal blocks that ignore distance. The first cluster you find is the network-wide common mode, and removing it is standard preprocessing; we remove it per kinematic sub-block, and every artifact of this deposit — the atlas, the player cards, the interactive explorers — shares one set of labels, so a block number means the same thing everywhere. **On the count, stated precisely because it is easy to get wrong and the atlas will show it:** across the 229 stations that survive to become lattice vertices, **fifteen chart labels appear** — 0 through 8 and 10 through 15. Those fifteen are what Figure 1 draws and what the common-mode step operates on. Three further labels (9, 16 and 17) exist in the clustering's numbering but carry no surviving station. Those three are not unused numbering. **They are real kinematic charts whose entire membership was removed by quality control**, and naming them matters because a chart that vanishes silently is indistinguishable from one that never existed. Chart 9 is BCNN and BCSC; chart 16 is OYLR; chart 17 is SKMA. All four stations were quarantined, none excluded. So the clustering found eighteen coherent groups, quality control emptied three of them, and fifteen survive to be drawn. Every figure in this paper is generated after that filter, which means those three charts appear nowhere — and that is a property of the quality gate, not of the ground. Each block is judged against its own reference.

That per-block discipline was paid for with a lesson. An early network-median subtraction manufactured a beautiful false anomaly: a southern Oregon family that appeared to anti-correlate with its neighbors at -0.87. The stations were innocent. The network-wide median is dominated by the north, so subtracting it from a southern group removes that

group's real shared motion and injects a spurious anti-correlation into what's left. The flags were reading the block-median-removal artifact, not the crust. The fix — judge each block against its own median — became load-bearing for everything after, and it planted the theme this paper keeps returning to: the thing you subtract is not a nuisance, it is a measurand you have chosen not to look at yet.

Stage two freezes the rules: a training era (weeks 500–1094 of the record, through 2020) for building and calibrating, and a vault — the held-out era, 2021 onward — that no instrument touches until the one opening described in §6. The southern segment forced one honest fork here: it has zero strict-quiet weeks (its slip rate at depth exceeds convergence; its episodes recur every eight to twenty-two months), so the working calibration tier excludes only major episodes. Hold that fact; it returns twice.

Stage three lays the lattice. Delaunay triangulation over the 229 surviving stations gives 441 triangles, every one oriented counterclockwise by construction (an unoriented mesh scrambles charge signs map-wide — a selftest catch, not a war story). Each station's week is one complex number: east and north residual as a small clock hand, a rotor. For each triangle each week, sum the wrapped phase differences around the loop; the result divided by 2π is an integer — the triangle's defect charge, +1, 0, or −1 — provided all three hands clear a one millimeter amplitude gate, because the phase of a near-zero vector is noise and the vortex core is rightly unreadable. Discrete Stokes is machine-verified on the mesh: sub-charges sum to their enclosing loop's charge, so every polygon's winding is the sum of the triangle charges it encloses, and 2.8 million candidate rings collapse into one 441-triangle computation plus addition.




Figure 1. The station and block atlas: all 229 surviving stations coloured by which of the fifteen populated kinematic charts they belong to (three further labels carry no surviving station), with the 46.11°N divide marked, the dipole stations and the inland-Oregon micro-block labelled by name, and the frozen witness pool ringed. This is the map every other figure in this paper is drawn on. The interactive block-explorer supplement in the deposit (BLOCK_EXPLORER_20260819_224258.html) draws the same map with every chart as a toggle, so a reader can turn blocks on and off by selection. (station_map v1.0, run 20260818_144812.)

One fence on scope, because a lattice invites it: an edge transmits exactly the one observable its connection is defined on — for GPS, one projection of the displacement field into a phase. Want strain too? That is a second connection, a second ledger, built and validated separately on the same graph. Multiple bundles over one lattice is legitimate and powerful, but each one is earned, not inherited.

The methodology this paper runs on — the noise-floor and pre-registration disciplines, the competence surface, the mod-n limit and its discrete-Stokes redemption, and the vocabulary tables of §8 — is documented in the companion *Loop Walking: A Field Guide to Knowing Whether You Found Anything* (concept DOI 10.5281/zenodo.21811588, which always resolves to the newest deposited version). The area-law measurement reported in §2 is the source of that guide's Rule 10 and of its answer to the loop-size question it had carried open since draft v4.

§2 — What the Training Era Showed

The first frozen scores broke the instrument's naive story and handed us a better one. Charge rates everywhere sit *below* the random-phase floor (measured with the tool's own function: 0.2509, against observed rates of 0.15–0.22) — the residual field is spatially correlated even in quiet, and *more* ordered during transients. North of the divide the paired test makes it a finding: 107 of 178 triangles dip during activity (sign-z +2.7), while the south stays null. Transients do not roil this field; they *order* it, suppressing the noise vortices. We call it coherence inversion, and it converts a charge-rate dip into a candidate transient marker — registered as a fork, scored later, never tuned.

Noise vortices are sign-symmetric; the map is not. Eighty-six triangles hold statistically one-signed standing charge, led by a Mendocino-corner triangle charged in 60% of its weeks — 2.4 times the random rate — at +27 against –166. And on one shared edge in the Olympic coastal group sits the object this program watched longest: OCEN-P398-PABH standing pure +120 of 120 charged weeks, adjacent OCEN-P397-P398 standing pure –111 of 111. A dipole, in the first instrument that could have seen one — either the Euler $\pm 1/\pm 1$ saddle hypothesis made flesh, or one bad station shared between two triangles. Four instruments then queued up to adjudicate (§3 tells that story); none convicted a station, and their failure measured something better: 35% of stations hold a *stable* phase offset against their neighborhood (population median concentration $R = 0.40$). The residual field carries standing spatial texture, map-wide and persistent — the below-random correlations are standing in time. The dipole is a knot where the frozen texture happens to wind.

Two more training results complete the picture. Area first: straddle triangles led the raw charge-rate table, but they are 1.8–2.3 times larger than their peers, and under area normalization the ranking inverts (density: north 0.307, south 0.182, straddle 0.046) — at the rate level this instrument is a phase-coherence-length meter, and the right comparisons are distance-matched. And the coarse scale: pointing the winding instrument at the subtracted common-mode series themselves — the blind spot made measurand — showed one block's common mode winding against the entire map at up to 3.6 cycles. The who-is-who roster names it: the Puget Sound-Salish Sea block, 21 stations centered at 48°N in the strait — the core of the Wrangellia ETS segment. The subtracted common mode was a nuisance we learned to measure; that inversion, promoted from fence to headline, is the paper's method in one sentence.

The area law: charge cancels at every scale we can measure

The dipole and the standing texture both raise the same question, one the loop-walking method carried open for four drafts of its Field Guide: **how big should a loop be?** The size law says the enclosed-curvature term shrinks faster than the path term, so small loops lose signal — but it does not say where to stop growing. The mechanism that would produce a best size is not noise accumulation; it is cancellation, a larger loop swallowing opposite-sign charges that net to zero. No model we had contained it. A triangle-charge lattice measures it directly, because the opposite signs are individually visible: grow connected triangle sets outward from forty seeds, record enclosed area, boundary edges, and net charge at every step — discrete Stokes doing the arithmetic — and compare against a null that permutes the charges among the triangles while preserving the exact charge multiset and the exact geometry, so the only thing destroyed is the spatial arrangement.

The ratio of real to shuffled RMS is 0.70 at every scored scale, across a factor of fifteen in enclosed area (7,295–110,758 km²: 0.703, 0.707, 0.697); growing the loops to 300 triangles and 273,766 km² gives 0.721, 0.713, 0.768 — still flat, still near 0.70. Not an optimum. Not monotone. Flat. The real arrangement cancels by the same factor whatever loop you draw — about thirty percent more than random placement, at every size. There is no characteristic scale in this channel on this substrate, and that sits beside the training-era finding that charge rates run below the random-phase expectation even in quiet: two independent statements that the field is arranged to cancel.

And the loop-scale question itself came back not evaluable — a specification, with its mechanism stated, rather than a shrug. Only three area bins carried enough null spread to support a ratio, and an optimum claim needs more: with three bins the argmax lands in the middle one time in three by chance, and coarser binning only merges bins into fewer. The blocker is that small loops enclose too little sum to form a ratio from, so what the question needs is a denser lattice — more cells, so more of the area range clears the noise threshold — a requirement for the instrument, stated as a measurement. (*Provenance: loop_area_law v1.5, run 20260813_213148 — pre-vault by timestamp — on the canonical triangle_charges.json, train era wk500–1094.*)

§3 — Hunting Rogues: Decouplers in Quiet, Couplers in Transients

“Loud isn’t equal to effective. That lesson the stars and quiet quasars taught me.”

The one-paragraph version

A GPS network is a team of witnesses, and some witnesses lie steadily. Standard quality control catches three kinds of bad testimony: sudden jumps, wild outliers, and flatlines (“this station hasn’t changed in a year — flag it”). But there is a fourth kind that passes all three checks: the station that drifts — slowly, smoothly, coherently — because slow coherent drift looks exactly like tectonics. We built a cheap way to catch the drifters. Lay a triangle lattice over the network, read each triangle’s phase circulation as an integer, and a drifting station betrays itself twice over: its triangles light up with persistently one-signed charge, and its phase winds against every quiet neighbor with the same sign, year after year. Deep-geodesy people know monument instability exists; what they mostly don’t have is a cheap lattice instrument that localizes a drifter geometrically by its triangle signatures and convicts it by winding consistency against quiet neighbors.

The smart-12 picture: rotors, phase locks, and lightning triangles

Picture every station’s weekly residual as a small clock hand (a rotor): the hand’s direction is the phase, its length is how far the station sits from where the region says it should be. Healthy neighbors phase-lock — their hands swing together, and any triangle drawn between three of them encloses no net rotation: charge zero. When something rotates inside a triangle — real ground circulation or one bad clock — the three hands can’t close, and the triangle carries a charge of +1 or −1. That is the lightning-triangle map: charge flashing across the mesh week by week.




Figure 2. The lightning-triangle map running, at week 919 of the training era. 441 oriented triangles; red is +1, blue is −1; persistence trails show where charge has recently been. 64 triangles are charged this week for a net of −4 and the dipole is labelled. The vault wall sits at week 1095 — everything right of it was computed blind. A still from the interactive supplement that ships with the deposit. (lightning_map v1.1, run 20260819_224304.)




Figure 3. The block explorer: the fifteen populated kinematic charts with their station counts, the 46.11°N divide, and click-to-identify. The second interactive supplement, all blocks enabled; the footer names the three charts emptied by quality control. (block_explorer v1.1, run 20260819_224258.)

A rogue station is a clock whose hand slowly turns relative to everyone around it. Its fingerprint is *consistency*: it winds the same direction against every neighbor, while the neighbors stay mutually quiet. A real patch of rotating crust marks only the triangles at its core; a rogue drags its whole fan.

The two-channel gallery

Rogues show up in two distinct channels, and the difference is the finding:

Decouplers in quiet. In calm periods these stations disconnect from their neighborhood — flagged at clustering time (14 quarantined, 2 excluded), kept on a watch list rather than deleted, because a kill without a reason is not acceptable and because of the second channel:

Couplers in transients. During real events, some watched rogues suddenly *join the chorus*. Quarantined WARM topped the entire north in the wk1360 margin-wide event (22.7, 4.8 σ); quarantined RYMD carried the first pulse of the wk1146 specimen; excluded BCHO joined the wk1360 chorus at 4.2 σ . Stations too unreliable to be lattice vertices still testify as witnesses — solid, liquid, porous earth suddenly gripping is a physical hypothesis, and the watch list is how it stays testable.

The winding census: how a conviction is earned

For every station, accumulate the wrapped week-to-week change of its phase relative to each lattice neighbor across the training era — the relative winding, in whole cycles. Then demand the rogue fingerprint: same sign against at least 90% of partners, median winding at least 0.7 cycles, four or more partners. The first full census (run 163117) flagged 26 convictions and 37 suspects out of 216 scored stations.

Then came the fences, because the first list was too good to trust:

The chart fence. Common-mode cleaning is done per kinematic sub-block, so a station whose lattice partners live in a *different* sub-block is being compared across two different subtractions — and real block-against-block rotation manufactures the rogue fingerprint at every chart boundary. The cure is to re-run the census using only same-sub-block partners, where the confound vanishes by construction, and to measure the blind spot directly: point the

winding instrument at the subtracted common-mode series themselves. That block-versus-block matrix (run 180621) showed one sub-block's common mode winding against the entire map at 0.8–3.6 cycles — and the who-is-who run (183234) names it: the Puget Sound–Salish Sea block, 21 stations spanning 46.2–50.1°N with its centroid in the strait at 48.05°N — the core of the Wrangellia ETS segment. Inter-block rotation measured, the subtracted structure promoted to the measurand; whether that winding is the northern ETS cycle itself read as block-level phase rotation is a registrable check (a ~425-day line on that block's common-mode phase), not a claim.

The geography fence. A station across a real kinematic boundary (Tofino, across the strait from the Washington coast group) winds consistently against long hull edges for honest tectonic reasons. Long-edge and divide-spanning partners earn a FENCED tag, never a conviction.

The roster (run 180621, the fenced census)

Survivors — convicted inside a single chart, where the confound cannot reach: TRND (winding 4.9× what block rotation predicts), P734 (4.6×), P416, SC02, P437, QMAR, P167, WAOL, PTRF, P406, ORTA. TFNO survives numerically but stays geography-fenced.

Cleared — the seam's own rotation read through a torn chart: SNOQ, SC00. Their fingerprint dissolved entirely inside one chart: formally CM-shadow.

The seam-dwellers — untriable, which is itself the finding: the loudest first-census convicts (P062 at 6.4 cycles, LINH at 5.1, P733, WAEN, REDM, P373) all have lattice fans that cross chart boundaries, so no single-chart trial exists for them. They live exactly where the seams are — and the who-is-who run shows why: the clustering had already isolated them. Four of them (P062, P733, P733, REDM) are their OWN four-station kinematic mini-block in the inland Oregon Cascades band (42.4–44.3°N), moving coherently together and winding together against the map — one rotating micro-block, not four independent rogues (fence: a shared agency fleet with common firmware could mimic this; the steps database decides). LINH, P325, MUIR, CHLW and SQMS are kinematic singletons; BCWS+WAEN and DCSO+NWE3 are isolated pairs. One number keeps P062 open either way: 6.4 cycles exceeds the largest block-against-block winding (3.6), so pure shadow cannot reach it even in principle. The micro-block reading also explains why the same-chart trial cannot exist for them: their chart is the anomaly, and a court convened inside it would be asking the defendants to judge each other.

What a station would have to be doing, physically

The mechanism menu, in rough order of prior: monument wander (a pillar settling in soft ground drifts millimeters per year — “monument noise” is established geodesy literature); local hydrology (aquifer drawdown moves the ground for real, but privately — one-station tectonics); multipath environment creep (tree growth, snow, new construction); and the human factor — equipment left in a wrong state, or known-bad stations kept running because replacement is not in anyone's budget. Some of these rogues may be old news inside the data-collection community, internally quarantined without ceremony. If so, this instrument found them independently from the public data alone — and the list is offered in that spirit: a cross-check, not an accusation.

Why standard QC misses them: jumps, outliers, and flatlines are the three standard tripwires, and a 0.3-cycle-per-year relative drift of a 2 mm residual trips none of them. The lattice does, because consistency-against-quiet-neighbors is a question no single-station diagnostic can ask.

The equipment objection, answered before it's raised

A careful reader will suspect equipment — clock-sync work rippling through local stations during the 2019 replacement wave, a recommissioned station blipping its neighbors, some zone-wide recalibration nudging everything toward agreement. We suspected it too, and audited it. Three answers, one structural. First: these positions come from precise point positioning — every station is solved independently against global orbits and clocks, with no network adjustment, so a site visit *cannot* propagate into a neighbor's solution; that coupling exists in RTK and classical network adjustment, not here. The only genuinely shared channels are the constellation, the regional atmosphere, and our own common-mode subtraction — which §2 turned from an assumption into a measurement. Second: “zone recalibration” is not a thing in continuous GNSS; position falls out of processing, not sensor tuning, and the fleet-wide worries that do exist (antenna calibration models, frame updates) are processing-era effects that the homogeneous IGS20 full reprocess applies uniformly — a 2019-localized model artifact cannot survive it. Third, the evidence that cuts hardest: the documented equipment swaps on the accused stations are staggered across 2015–2019, while the coherent winding persists through every one of them. Equipment artifacts flinch at swaps; this ground does not. And the one station whose jump the step database *does* explain — P397, antenna and receiver swapped eighteen days after the GPS week-rollover of 2019-04-06, its jump dated by our instrument to within two days of the visit — demonstrates the dating accuracy and removes itself from the mystery pile: one story, two channels, documented swaps beside silent firmware glitches that leave no paper trail.

What the rogue hunt taught us about the instrument itself

Chasing the loudest rogue candidate — a pure $\pm 120/111$ standing dipole on one shared edge — through four instruments in sequence —

- steady rotation: acquitted;
- phase jumps: decoupled from the charges;
- burstiness against a measured null band: nothing above null;
- static bias: two stations anchor a stable offset, two hold none —

left no station convicted and something better discovered: the residual field's phase pattern is *standing* — persistent spatial texture, map-wide. The rogue instruments, failing to convict, measured the field. The dipole is a knot where the frozen texture happens to wind.

And one honest limit, documented with its redemption: a triangle reads winding mod 3 (a phase step of 240° per edge — a true +2 — reads -1; 360° per edge — a true +3 — reads 0). Three samples around a circle cannot tell a wheel spinning forward fast from one spinning backward slow. The anti-aliasing is discrete Stokes: sum the triangles, and a 200-gon reads mod 200. Full statement in §8.

How to check us, and how to run this on your own network

Everything here regenerates from public data and short pure-stdlib scripts shipped with the deposit: the census, the fences, the block-matrix, and the null band are all one double-click each, and every output file carries the tool's name and its run stamp. The deposit also carries the two interactive supplements (the lightning map and the block explorer), the canonical charge map beside the repaired-grouping map of §7.2 under its own name, and the full tool roster with checksums in Appendix C. The portable recipe is three steps: build the lattice, run the winding census with the same-chart fence, and read the block-versus-block matrix before believing any boundary conviction. For the 2019 jump cluster we flagged (seven stations' career-largest phase jumps in one six-month window), the receiver-side GPS week-number rollover of 2019-04-06 is the surviving candidate, and the closing test is public: the step database entries near each jump date.

§4 — The Silent Traveling Front

A word this paper leans on from here forward: an *onset* is the start week of a catalogued episode on this margin — the tremor catalogue's own episode boundary, or a catalogued earthquake — and the *onset channel* is our detector for stations firing together, built from the acceleration centroid of the firing pattern. The channel is not a forecaster and does not decide what an episode is; it watches how stations behave around dates the catalogue supplies.

Twice in the training record, the onset channel caught something moving that made no sound. In the weeks around wk1146 and again at wk1151, the same family of stations fired in sequence — two pulses, tremor-silent, with the acceleration centroid of the firing pattern migrating north along the margin at roughly 2-5 kilometers per day. That speed is the signature that matters: slow-slip fronts propagate at kilometers per day; every instrumental artifact we can name (processing waves, constellation effects, weather systems) moves orders of magnitude faster or not at all. A silent signal at slow-slip speed, read purely from GPS geometry, is the channel doing the one thing it was built to do: report motion without needing the ground to shout.

The stated methodological claim, with its fence: the acceleration centroid as a propagation tracker for silent transients is, as far as we know, not a standard tool. That's the contribution, and it is cheap enough that anyone can check it. The fence is that $n = 1$ (arguably 2): one traveling specimen family, one training era, no cycle-level statistics — which is why the claim ships as method plus specimen, not as physics. The contrast specimen sharpens it: the wk1360 margin-wide event lit 91 stations in a single week with no buildup and no migration — a different animal entirely, coherent and sudden where the front was sequential and slow. Two specimens, two morphologies, one instrument that distinguishes them. And the vault opening added an echo it could not have been tuned to produce: the Class-B charge-channel description at the wk1146 window (§6) shows a twenty-triangle cluster holding charge along the same Salish corridor, containing two stations of the front's own carrier family, peaking in the front's two pulse windows.




Figure 4. The lattice mid-episode in the held-out era, week 1107, with persistence trails on: 30 charged triangles, largest cluster 4. **This frame was selected by eye for illustration and demonstrates nothing** — no criterion picked it, and a still cannot establish propagation. It shows what the instrument displays during an episode, not evidence of a front. (lightning_map v1.1, run 20260819_224304.)

The open community challenge lives here, as this section’s closing gesture — full ledger, tools, and both specimens ship CC-BY with an explicit invitation to explain them better than we can. The invitation stands as drafted: here are two specimens, every tool that found them, and every fence we could build; if you can break them, the breaking is a contribution and we want to read it.

§5 — The Witness Pool

If some stations lie, which ones tell the truth — and can you know in advance? The superstar scorecard asked it with six criteria, each named and thresholded before scoring: census-clear in both the full lattice and its own chart; winding-quiet; jitter inside the measured CLEAR null band; near-full coverage; alive above the amplitude gate; and present in the fan during both training-era event windows. The design law was consistent-is-not-loud — no criterion measures amplitude — and the trap was a survivorship null computed from the criteria’s own base rates.

The trap snapped shut on our own hypothesis, which is the design working. Two stations passed all six; chance expects 1.5 from 97 fully measurable — and the criteria correlate, so even that understates chance. The dozen-superstars idea died on training data, where dying is free. What lived is better: reliability is a *graded spectrum* — 26 stations at five-of-six, structure everywhere — and half the network (101 of 228) cannot even be cross-examined in its own chart, a measurability fact the binary badge had been hiding. So the pool rule was registered before its own selection run ever executed: pass/fail/unmeasurable per criterion, eligibility at $n_{\text{pass}} \geq 5$, ranked by passes then measurability then winding-quietness, capped at thirty, surprises ship. The frozen pool came out at 28 names (run 212659) and became, by that act, the vault’s witness list.

The pool’s shape is itself a finding. Blocks 11, 12, and 1 supply nearly all of it; the Puget-Salish block — the coarse instrument’s loudest winder — contributes zero witnesses, the fine and coarse instruments agreeing without being asked. The Mendocino honest family (P158, P161, P162) sits quietly verifiable at the feet of the most-charged standing triangles on the map: honest stations, reporting extraordinary triangles. And P398 — the dipole’s acquitted star — made the pool at five passes with one criterion unmeasurable, a witness whose own chart cannot cross-examine him.




Figure 5. The frozen 28-station witness pool, green-ringed, on the block explorer with only the contributing blocks enabled. Blocks 11, 12 and 1 supply nearly all of it; the Puget-Salish block — the coarse instrument’s loudest winder — contributes none. (block_explorer v1.1, run 20260819_224258.)

Every station’s full accounting, pool and field alike, is Appendix A.

§6 — Opening the Vault

Before any instrument touched the held-out era (wk1095–1383, roughly 2021-01 through 2026-06), we registered five claims with their metrics, nulls, and thresholds, froze every tool by checksum, and committed to one run whose first rendered result would be final. The vault opened once, on 2026-08-14. The results: the frozen 28-station witness pool **validated** against its null (V1); the standing-texture channel **validated**, with the OCEN dipole surviving by name with both signs intact (V4); two claims (V2, V3) were killed before measurement by a **design error of ours** — their registered nulls required quiet weeks that our own earlier measurements had already shown do not exist on this margin — and were re-run as labeled post-hoc tests so the questions do not go unanswered; and the Class-B storytelling channel (V5) recorded sustained charge clusters at both pre-known event windows, including the dipole itself sitting inside the 2026 loop. The pipeline integrity gate reproduced the entire training era exactly, 441 of 441 triangles, before a single held-out week was read.




Figure 6. The instrument standing at the vault wall, week 1095, with the same layers as the training-era frame so the two can be read side by side. The era readout shows HELD-OUT (vault); 25 triangles are charged for a net of +1, against 64 and -4 in training. Everything from this week forward was computed blind and opened once. (lightning_map v1.0, run 20260814_025141.)

The one-shot design, in brief

The training era built and calibrated everything; the held-out era was never touched by the charge, winding, or witness instruments. VAULT_PROTOCOL v0.1 registered each claim's blindness class in advance: Class A (fully blind — can confirm), Class B (locations pre-known from our own onset channel — description only, no pass/fail), Class SPENT (the onset channel itself, which consumed the full era during earlier calibration — exploration only). One double-click, first rendered result final, all outcomes shipped with their labels. The run's own first act was to re-earn the opening: recompute the full training era from raw cache and demand an exact match to the canonical record. It matched, 441/441.

V1 — Witness-pool stability: VALIDATED (Class A)

The 28-station pool was selected on training data under a rule frozen before its own selection run, and named in the protocol before the vault opened. Held-out re-measurement: 8 of 28 retained the reliability bar ($n_{\text{pass}} \geq 5$ of 6 criteria, C3 band re-measured from held-out CLEAR stations at 0.27-0.30), while 1000 random 28-station pools from the remainder retained a median of 3 ($p_{95} = 6$). $S = 0.29 > 0.21$. **The pool carries real, transferable information about which stations to trust.**

The retained eight: P324, P337, P349, P407, P408, P410, P793, UCLU — with P410 and UCLU scoring 6-of-6 in an era they had never seen. Honest movement in both directions: GOLY, one of only two training-era 6/6 stations, fell to 3; P736, WAMC, and P376 fell to 1. Reliability is graded and it drifts; $S = 0.29$ is its measured five-year transfer coefficient, and the drift itself is data for the dynamic-reliability program of the next era.

V4 — Standing-texture persistence: VALIDATED (Class A)

The training era found 86 triangles holding statistically one-signed (“standing”) charge. Of these, 28 remained measurable held-out (≥ 30 charged weeks under the unchanged gate — the held-out era is half the training era’s length, so reduced measurability is expected). Sixteen stood again; **fifteen of the sixteen kept the same sign** (binomial $p = 2.6 \times 10^{-4}$ against sign-symmetric chance; persistence $0.54 > 0.5$). The standing texture is not a training-era artifact: it is a persistent spatial configuration of the deformation field, confirmed blind.

And the single most-watched object in the program — registered by name with no threshold, “watched by everyone, claimed by no one” — survived:

OCEN-P398-PABH: train +1 → held-out +1, STANDING. OCEN-P397-P398: train −1 → held-out −1, STANDING.

The adjacent pure-opposite dipole pair held both its signs through five years of unseen data. Four training-era instruments had already shown it is not rotation, not jumps, not bursts — a standing phase arrangement. Whatever it is, frozen strain gradient or stubbornly static bias, it does not blink.

V2 and V3 — killed by our design error, not by the data

(four lines, per the project’s own standard, then the mechanism)

V2 and V3 report NO VERDICT: their registered nulls required event-free held-out weeks, and the compiled catalogs cover 289 of 289 — a fact our own stage-2 freeze had already measured (southern strict-quiet = zero) and that we failed to carry into the protocol’s null design. This was an error, ours, knowable in advance; both questions are answered post-hoc below rather than left unmeasured.

The mechanism, measured. With every catalog episode padded ± 2 weeks and every in-latitude $M \geq 6$ padded ± 4 , the 154 compiled windows blanket the entire held-out era. The southern (Klamath) segment alone — 84 of the 147 ETS windows — covers 99% of it, which is the segment’s known physics: the shortest recurrence on the margin, slip rates at depth exceeding convergence. On this margin, at weekly resolution, **there is no quiet time**; “quiet” in the south has always meant departure-from-self, not absence of events. That is itself a finding, and it dictates the corrected design: measure departures from the margin’s own busy baseline, not distance from a silence that does not exist. To be explicit about what was *not* lost: all 229 stations loaded, all 441 triangles (214 of them southern) reproduced and extended through the held-out era — the data is whole; the comparison design was wrong.

Integrity note for the record: no V2 or V3 statistic was ever computed and withheld. The first attempt crashed on the empty null set before any such number existed (the crash report contains only a traceback), and the authorized fix predetermined the no-verdict outcome before the re-run. The timeline is in the run artifacts.

Phase 2 — the same questions, run down after the fact

[POST-HOC — label carried in every header. These numbers cannot claim pre-registration; they ship verbatim, favorable or not.]

Design (the departure-from-self null): the corrected comparison whitens against the margin’s own baseline instead of demanding silence. **V2-posthoc:** each catalog-selected event window’s peak edge-connected charged-cluster size versus N random windows of identical length placed uniformly anywhere in the held-out era — typical margin weeks, background events included — paired sign test, direction as originally registered. **V3-posthoc:** per-triangle charge rate inside major episodes versus that triangle’s own rate in the busy background, north block claim-shaped as originally registered (dip), south computed alongside.

V2-posthoc (run 20260814_010907, seed 20260814): NOT SUPPORTED under either null — 67/154 windows beat the anywhere-null medians ($p = 0.95$); 47/124 beat the majors-excluded null ($p = 0.997$, and significantly BELOW half, low-tail $p = 0.0045$: event windows carry smaller peak clusters than the busy background). Full results, per-event table, and the inversion-direction observation: Addendum. V3-posthoc (same run): the registered north dip does NOT transfer — north 78 dip / 97 rise, sign- z -1.44 (rise-leaning, below any bar); south 66/55, $z +1.00$, computed beside. Full results: Addendum.

The corrected nulls are simultaneously the registered design for the next sealed era (wk1384 onward, accruing weekly), where they enter through the front door.

V5 — the charge channel at the known events (Class B: description only)

Both windows were located by our own onset channel during training-era work; selection is contaminated and no pass/fail exists here by construction. What the lattice did there, verbatim:

wk1146-front window (wk1120-1177): a 20-triangle edge-connected cluster along the Salish Sea / southern Vancouver Island corridor (AL2H, ALBH, CLRS, PTRF, PGC5, BAMF, UCLU families) — containing CLRS and PTRF, two stations of the silent traveling front’s carrier family — holding charge through the window with peaks of 6 simultaneous charged members at wk1131 and wk1144, spanning the front’s two pulses.

wk1360-margin window (wk1350-1370): a 19-triangle Olympic Peninsula / Puget cluster (CUSH, BILS, DEEJ, PABH, the P400 family) sustaining 5–7 charged members for fifteen consecutive weeks, peaking at 7 in wk1364 — **with the OCEN dipole pair itself (OCEN-P398-PABH and P398-WAMS-PABH) sitting inside the loop.**

Description, not claim — but the lattice was telling stories precisely where the onset channel said to look, and the storytelling-loop methodology these windows exercise is exactly what §7’s post-hoc battery scores in this paper, and what the next era’s registered V2 will score again, corrected.

Scoreboard

Of four Class-A claims, the two that our design allowed to be measured both validated — the geometry channel (who to trust, and what stands) transfers blind across five years. The two event-contrast claims died of our null-design error, are answered post-hoc in the Addendum, and re-enter the next era corrected. The Class-B channel recorded rich structure at both known events. The vault is spent; every idea born during the opening belongs to the next era.

§7 — What the Post-Hoc Answers Add

The Addendum carries the full tables; the short version belongs in the narrative. Measured the only way this margin permits — against its own busy baseline — the event-contrast channel shows no positive transfer: catalog windows do not out-charge the background, and the northern charge-rate dip does not reappear. One tail is significant and points the *other* way: under the majors-excluded null, event windows carry smaller peak charge clusters than background more often than chance allows ($p = 0.0045$). That is the coherence-inversion direction — the training era’s frozen finding that transients order the field — knocking from the far side of the wall. It is recorded, fenced twice over as post-hoc, and it becomes a front-door registered claim (“V2-inverted”: charge clusters dim during catalog events) for the next era. Nothing in the post-hoc battery touches V1 or V4, which validated under registered, feasible nulls.

7.1 Does the ordering lead? (POST-HOC, run 20260819_171805)

The program has never asked whether the lattice’s ordering signature *leads* a catalog onset or merely accompanies it. This section asks it, on the spent record only, under a design frozen before the data was touched.

The design fixed everything in advance: leading windows running from onset minus k weeks to the week before onset, with k pre-stated as two, four, eight and twelve; major ETS episodes plus in-latitude $M \geq 6$ earthquakes; both departure-from-self nulls printed side by side; and events, never windows, as the independent unit, so the nested ladder is never counted as four separate pieces of evidence. The integrity gate recomputed the entire training era from cache and matched the canonical record on all 441 triangles before a single scored week was read.

A correction belongs here, before any result, because it changes what this section is entitled to claim. An earlier build of this analysis reported that 110 onsets reduced to 25, and said nothing about what kind of event each was. Behind that formatting choice sat two defects. Same-week collisions were resolved by keeping whichever event had been added to the list first, and ETS episodes were added before earthquakes — so an earthquake sharing a week with an episode was discarded silently, with the precise object (a point with a timestamp) losing to the vague one (a segment-wide episode dated to about a week). And a single twelve-week separation fence was applied to every window length at once, which on a margin where southern ETS recurs every few months is close to a total exclusion of earthquakes. **Of 22 earthquakes at $M \geq 6$ in the spent era, three were scored. The largest event on the margin, an M7.0 on 2024-12-05, was never asked the question and appeared in no output.** Nothing was computed wrongly and no window was contaminated — the fence still saw every week that carried a quake, because the collision discarded a duplicate on a week and never the week itself — but the analysis could not say what it had and had not tested. The rebuilt tool merges collisions instead of dropping them, decides cleanliness separately for each k , and scores every event whether contaminated or not. **This section reports the rebuilt run.**

The spent record carries 110 event-weeks: 92 holding ETS episodes alone and 18 holding earthquakes, with 22 quakes on those 18 weeks because two aftershock sequences put three each on a single week. Six weeks carry more than one event and are now merged rather than thinned. Cleanliness is decided per window: **89 of 110 event-weeks are clean at $k = 2$, 70 at $k = 4$, 42 at $k = 8$, and 25 at $k = 12$.** The old single fence returned 25 for every window length, punishing the short windows for a contamination only the long ones suffered.

Every $M \geq 6$ event now appears in a printed ledger with its week, whether it is clean at each k , and — where it is not — which events share its window and how many weeks before. The M7.0 is scored, clean at $k = 2$ and $k = 4$, and contaminated at the longer windows by an ETS episode eleven weeks prior and an M6.0 five weeks prior. That is a statement the earlier build could not make in either direction.

The answer is no, and it does not change when the earthquakes are added back. Across the sign tests — three readings, four lead times, four spatial shapes, two null designs, and now two scored sets — **not one cell falls below a p-value of 0.14**, the most extreme being 0.1449. The pre-stated dimming direction did not appear. If anything the peak reading leans the other way, with p-values as high as 0.999 in the anti-dimming direction, which is to say event windows carried *larger* charge clusters than background more often than chance allows. That is the same coherence-inversion direction §7's opening paragraph already records from the vault's post-hoc battery, arriving again from a different instrument, and it is registered as a front-door claim for the next era rather than argued here. The ternary reading returned heavily cresting throughout, which is what a background-like result looks like when the band is the null's own central 68%.




Figure 7. The superposed epoch for the uniform mode — the answer at a glance. Each trace is the ordering statistic averaged across scored events, aligned on onset at zero; the shaded band is the matched-length null. A leading signature would show as a departure before zero. There is none. (PRECURSOR_POSTHOC v2.1, run 20260819_171805.)

Per-event leading-window strips — the ledger's weekly step

POST-HOC · run 20260819_171805 · shaded = the k=12 leading window · dashed = onset · raw, unsmoothed · showing 24 of 110 events




Figure 8. Per-event strips, one panel per scored onset, so a reader can see that the flat epoch is flat everywhere rather than an average cancelling two opposite behaviours. The earliest panels begin at the record's own edge — the record opens in 2009, so the first onsets carry no leading history to draw, which is the record starting, not data missing. (Same run.)

One number bounds what this null means, and it matters more than the p-values. The four margin-wide shapes the analysis reads — the whole margin together, north against south, ends against middle, and one finer shape — hold between them **3.1% of the weekly ordering**. A spatially structureless field would place about 5.7% into four of seventy available directions by geometry alone. Cascadia's weekly ordering therefore carries *less* margin-scale coherence than noise would. So the honest statement is narrower than it first appears: **the ordering does not lead, in the four smoothest shapes this margin has, and those shapes are three percent of what the lattice does each week.** The remaining ninety-seven percent is finer-grained than any of them and was not examined by this design. A flat epoch was pre-agreed to ship at full weight, and it does.

7.2 Does it survive its own preprocessing? (POST-HOC sensitivity)

A defect in the common-mode step was found after the analysis was complete, and because it is the kind of thing a reader is entitled to worry about, it is reported here with the measurement that bounds it rather than as a footnote.

What the defect is. §1 says common mode is removed per kinematic sub-block. The code selects a station's group with an expression that treats a missing label as a signal to fall back to the coarse superblock — and the chart labels are stored as integers, so **chart 0 was read as “no label”** and its 48 stations never formed a group. They fell through to the superblock and split at the 46.11°N line into 37 northern and 11 southern stations, each half referenced against a different median. That is also where the count “sixteen groups” in the run logs comes from: fourteen real charts plus two fallback buckets, not sixteen charts.

Why it is not a post-hoc choice. The expression cannot see outcomes. It applied identically in the training era and in the vault, from before any result existed, so it could not have leaked held-out information. **The vault remains a valid one-shot test of the instrument as it actually was.** This is a specification defect — the paper described a procedure slightly different from the one that ran — and the correct response is to measure it.

The measured effect. Recomputing every charge with the grouping repaired and diffing against the canonical map: **11.16% of triangle-weeks change, 14,463 of 129,556**, rising to **27.97%** among triangles that touch chart 0 and 11.61% within the training era alone. That is not negligible, and it is larger than a geometric argument predicted — because a chart with fewer than five members does not use its own median but reaches for the whole superblock, and the cleaning runs in place, so changing the group list changes what those small charts read. The effect propagates beyond chart 0's own triangles.

Every headline conclusion holds under both groupings. The full analysis was re-run against a repaired charge map, and the results sit beside the canonical ones. One row needs its term in advance: a *tattler* is a station that gives away a coming onset — one whose winding, recovered from its neighbours while the station itself is blinded, reads non-background in the weeks before an onset. The fraction below is how often that happens across scored stations, and the background line under it is what construction alone produces.

reading	canonical	repaired grouping
tattler non-background fraction (median)	0.26	0.25
— against a background of	0.32	0.32
rhythm accel / decel / steady	7 / 25 / 68	7 / 24 / 69
rate-of-rate, onsets minus null	0.000	0.000
ring-walk lags, ring 1 / ring 2	25 / 19	25 / 19
leading-window ladder	null throughout	null throughout
the V4 dipole, both triangles	signs hold	signs hold

The dipole is worth naming because it is this paper's flagship blind confirmation and both its triangles carry chart-0 vertices, putting them inside the most-affected zone. Their net charges move — **+120 to +132 and -111 to -127** — and **both keep their sign**. The dipole does not depend on the defect.

One cell deserves stating rather than burying. Across the whole sensitivity ladder the minimum p-value is **0.0045**, in a single FULL-set NULL-B cell whose NULL-A twin reads **0.4617**. A result that appears under one null design and vanishes under its twin is a flicker, not a finding; corrected across the family it reads **0.14**. It is recorded here so that nobody discovers it later and wonders why it went unmentioned.

Everything in this subsection computed from the repaired map is post-hoc by construction and is labelled so wherever it is quoted. The vault was opened once, against the pipeline as it actually was; that blindness belongs to the as-built recipe and is not transferable to numbers computed afterwards. The repaired grouping is the next era's instrument, not a retroactive replacement for this one. The canonical charge map was never modified, and the repaired map ships beside it under its own name.

(Canonical run 20260819_171805; sensitivity runs 20260819_213054, 213058, 213103 and 221021.)

§8 — Limits and Vocabulary

The instrument’s honest ceilings, stated by the people who built it. A triangle reads winding mod 3: 240° per edge — a true +2 — reads -1; 360° per edge — a true +3 — reads 0; three samples around a circle cannot tell a wheel spinning forward fast from one spinning backward slow. The redemption is discrete Stokes — an n -gon composed of triangles reads mod n , so the big composite loops hold the big windings the small ones alias — and the physical fence is that a true [2] core inside one ~ 40 km triangle would need the residual vector rotating twice along the perimeter, implausible at our amplitudes: a documented ceiling, not a live wound. The amplitude gate makes the vortex core itself unreadable by design — the phase of a near-zero vector is noise — so the instrument reports the wall of the vortex, never its eye. Seam-dwelling stations are untriable by the same-chart census, and that is a statement about method, not innocence: they need a cross-chart instrument, and until one exists their verdict is UNRESOLVABLE, a vocabulary word this project uses in place of a guess. The 350–352-day band is the GPS draconitic year and is treated as an artifact channel wherever it appears, regardless of how exciting it looks. And the phrase “defect charge” means the wrapped-phase circulation integer of this paper, nothing grander: the full glossary of terms that sound alike across disciplines, false friends included, is Appendix B.

The sum is not bounded by the walk’s ceiling — measured. The mod-3 redemption now has a number, and it runs in the direction that is easy to miss: an enclosed sum is not bounded by its own boundary’s mod- n ceiling. On a synthetic coherent field, 120 summed triangles return 120 while the enclosing 34-edge boundary could only ever have represented ± 17 — walking the boundary wraps the number away; summing the cells inside recovers it. On the real map the effect is present but rare: one of 4,800 grown loops carries a sum exceeding its boundary’s ceiling, against 92% on the planted-coherent synthetic — itself a statement that the real field is far less coherent than a planted patch, which is what the below-random charge rates already said. The consequence for anyone building on this: a large loop is better treated as a query over small ones than as a walk; the n -gon is a sum, not a path, and the sum is the more informative of the two. One fence, stated because it is tempting: the largest integer any loop held grows with area (1 through 7 across the bins, 9 at deeper growth) and tracks the mod- n ceiling closely — that is expected, not a finding, because boundary count and the extreme of a random-sign sum both grow like the square root of area; only a null-relative spread statistic separates structure from scaling.

Two instrument lessons the area measurement cost, recorded beside the read-contract rule they extend. First, the statistic was wrong before the answer was: the first reader reported the contrast as a median and every bin came back unscorable — while the maximum went 1, 3, 5, 8, 10 across the same loop sizes. A triangle charge is a discrete, zero-centred integer, and the median of such a quantity reads its quantization, not its spread; this program’s house style is robust statistics, and robustness is the wrong instinct here. RMS grows like the square root of the loop count under random signs and linearly under coherence — an elevenfold separation, on a planted fixture, that the median could not see at all. Second, charge records merge by label, never by position: each triangle’s charges are keyed by week and defined only where the amplitude gate passed, so a positional stack concentrates every triangle’s first-defined week at index zero and manufactures coherence out of sparsity. Measured cost of that error, before and after realigning by label: the largest integer the lattice appeared to hold fell from 12 to 7, and the fraction of loops exceeding their boundary ceiling fell from 27% to 0.02%. The wrong version looked better. The realignment validated itself — the union of all 441 label sets reconstructed the frozen training era to the week — and both lessons are the pipeline’s version of a wrong accession: a sum across mismatched indices makes every internal check come out honestly wrong together.

Count the independent draws, not the rows. The grown loops above are nested — every larger set contains the smaller, about forty-four per seed — and are not independent samples. Scored across loops, a contrast peak cleared its neighbours comfortably and an early version printed “interior optimum present”; scored across the forty seeds, the actual independent unit, the same gaps fell inside the error bars and the verdict was flat. A factor of about six and a half in significance, from one wrong denominator. The same rule binds everywhere this paper computes a statistic over overlapping windows, fans sharing a vertex, or triangles sharing an edge: before writing any error bar, name the smallest set of independent draws the samples are built from, and compute the spread across those. The disease is not ours alone — the MeerKAT Pulsar Timing Array’s first gravitational-wave search reports that different noise assumptions yield either compelling Hellings-Downs correlations at $3\text{--}3.4\sigma$ or insignificant ones, apparently because its highest-precision pulsars sit close together on the sky and dominate the statistic (Miles et al., arXiv:2412.01153).

Two prior arts own neighboring ground, and we cite them here before a reviewer does it for us. The loop over a sensor array is sixty-eight years old: the closure phase (Jennison 1958, MNRAS 118, 276–284) sums visibility phases around a triangle of antennas so that the per-antenna error terms cancel — the Byzantine-sensor problem solved topologically in 1958, foundational to VLBI straight through to the black-hole images, and its bookkeeping already knows discrete Stokes: N stations yield of order N^3 triangle sums with only $(N-1)(N-2)/2$ independent among them. Interferometry has even arrived at the gauge language independently, framing closure quantities as the invariants of an Abelian gauge theory (Thyagarajan, Nityananda & Samuel 2022, Phys. Rev. D 105, 043019). What that literature does not do is the census’s job: closure *cancels* a bad antenna’s error without ever naming the antenna; the lattice names the station and convicts it. And the charge object itself carries a prior name: in InSAR phase unwrapping (Goldstein,

Zebker & Werner 1988, Radio Science 23(4), 713–720), phase sums around closed paths define *residues* of charge ± 1 , branch cuts are laid between them as barriers to integration, and most residues arrive in $+/-$ pairs — our dipole, under its existing name, complete with dipole-remover preprocessors in the MRI unwrapping literature. The positioning is the inversion this corpus keeps meeting: that field computes residues in order to remove them. Their obstruction is our observable; their nuisance is our signal. It is the third instance of the conjugate-moment structure (muon scattering, AlphaKnot, now InSAR), which is beginning to look like a law about this instrument rather than a coincidence. The full credit ledger lives in Appendix B.

§9 — The Next Era

The record keeps growing: wk1384 onward is a fresh sealed era that no instrument has touched, accruing weekly. Its registered slate is drawn from this paper’s endings. The corrected event-contrast tests go in through the front door — V2 under the departure-from-self nulls, the V2-inverted candidate (clusters dim during events) with its own threshold, V3 as majors-versus-background — and this time the null designs were checked against the project’s own coverage measurements before registration, which is the after-action rule this paper’s one confessed error purchased. Dynamic reliability graduates from pilot to claim: witness-phase state-stability with the baseline-matched per-station nulls the pilot’s confound taught us to build. The Puget-Salish block’s common-mode phase gets its registered ~ 425 -day metronome check. And one audit is designed and waiting: time reversal. Walking a ring backward flips winding by construction — an identity, good for auditing code, silent about physics. But playing the *record* backward is a genuine test: real tectonic loading is odd under time reversal, so run the reversed series through the whole pipeline, add the forward and reversed ledgers, and genuine winding cancels itself — whatever survives the sum is even under T, and “time-invariant dipole” is exactly the right name for that residue, because it is the artifact channel, isolated by symmetry instead of argument.

Further out, labeled speculation because it is one: each composite ring has a loading rate — a pitch — and pairwise rate ratios between neighboring rings are intervals. Whether those ratios sit on rational plateaus or drift is the mode-locking question, circle maps and Arnold tongues on crust, and if coherent blocks are real, the spectral families and the mode-locked families should be the same families found two ways. The fabric may play. That is next-era music, written down here so it can be tested there.

Safety and data currency

This paper describes instrument characterization on archival data; it makes no forecast. The onset channel cannot speak to the present in any case, and the reason is in its own construction: it smooths over four weeks before differencing and then zeroes the edges, so that a boundary artifact cannot masquerade as an onset. The most recent month of any pull is therefore blind by design, and the freshest verdict this channel can give concerns a window already about a month old. With that limit stated: the last check of the data edge was made on the 2026-08-11 pull, and every onset spike in the scoreable window carried a catalogue label — an ETS episode or a catalogued earthquake — with nothing unexplained. Should such a check ever disagree, the protocol-not-forecast framing and the appropriate notification path take precedence over publication schedule.

§10 — Addendum and Epilogue: New, Weird Ground

You are standing on the low ground now. Part I rode downhill: a registered question, a frozen vault, a one-shot answer, and a precursor campaign that ended flat and honestly. What follows is the wide, strange country between the descent and the climb — the post-hoc exploration that taught us the reframe, which the Interlude then names and Part II makes falsifiable.

Now for something completely different — the line is Monty Python’s, and borrowed on purpose, because what follows really is a different programme. Usually the addendum sits at the end, buried — it’s where people stick charts, graphs, bits of findings and post-hoc musings. Bringing it forward is unconventional, and so am I. This project was my very first triangle-lattice project as a differential-geometry engineer and Loop Walker, and I needed to push the lattice to its limit. I wanted to BREAK it — that’s how my brain works, and the lattice is a tool of the Loop Walker: I wanted to understand when and how it would help me hack substrates. Many of the most interesting, wild and amazing items were found post hoc, after the official project was finishing. Why include them? Because many folks will struggle with some of these concepts and I wanted to honestly over-share. Because this is new and exciting ground. Because I will never be as good at this as someone who does it for a living. But the post-hoc addendum is where the fun stuff lives — the reason we wake up in the morning, ask questions, and do science to understand the world. I hope you enjoy it.

Everything in this addendum is post-hoc exploration on the spent record. None of it was pre-registered, none of it is a claim about what will happen next, and several results below raise more questions than they settle. That is why they are included. We are reporting what the instrument did, not explaining it, and where a mechanism is offered it is labelled as a guess we have not tested.

10.1 The rhythm: this ground does not tick, it bursts

Every statistic to this point fixed the week and asked how much winding happened inside it. It is also possible to fix the *winding* and ask how long it took. A charged week is one completed turn, so the gap from one charged week to the next is exactly one full wind's worth of time. The inversion costs nothing, uses the same charge map, and produces integers.

Across the spent record the lattice completed **20,087 windings** over 421 triangles. The median gap between one completed wind and the next is two weeks, and in 44% of cases the gap is a single week.

The two rhythms are wildly asymmetric, and this is the most striking descriptive result the program has produced. Separating the gaps by whether the next winding went the *same* way or the *opposite* way: winding again in the same direction takes a median of **two weeks** across 17,515 gaps; reversing takes a median of **eighteen weeks** across 2,572. Nine times longer to turn around, and nearly seven times more same-way gaps than reversals. Once this ground starts winding in a direction it strongly prefers to keep going, and undoing itself is both rare and slow.

And it has no metronome anywhere on it. Measuring the spread of each triangle's gaps against their mean — near zero a clock, near one memoryless, above one bursts separated by silences — not one of 371 triangles came in below 0.75. Nine percent sit in the memoryless band; **91% are bursty**, median 1.81, one triangle at 6.45. The lattice winds in episodes.




Figure 9. How long one full wind takes: weeks between one completed winding and the next, whole lattice. (WINDING_RHYTHM v2.3, run 20260819_174255.)

Does the ground tick or burst?

spread of the gaps over their mean, one value per triangle — 0 is a metronome, 1 is memoryless, above 1 is bursty
WINDING_RHYTHM v2.3 · run 20260819_174255 · integers only, no amplitudes




Figure 10. Does the ground tick or burst — the spread of the gaps over their mean, one value per triangle, with the memoryless line marked. Nothing falls in the metronome region. (Same run.)

Those two findings are one structure seen twice: episodes are runs of same-signed weeks, and a run of same-signed weeks is both a short same-way gap and a long silence afterwards. On a margin whose defining behaviour is episodic tremor that is a sensible thing for the instrument to be seeing — though we have not demonstrated the two are the same episodes and do not claim it.

Does the rhythm change before an onset? No, in either scored set. Taking the last four gaps completed before each onset against that triangle's own habit, with the verdict cut at the null's own sixteenth and eighty-fourth percentiles so the test is calibrated by construction: the clean set reads 7.2% accelerating, 24.6% decelerating, **68.2% steady** against 68% expected; the full set including every contaminated event reads 6.3 / 24.4 / **69.3**. The rate-of-rate — whether gaps shorten monotonically toward onset — reads +0.200 before onsets and +0.200 at randomly placed weeks on the same triangles. The difference is exactly zero in both sets.

A caution we paid for. The first version of this analysis reported 1% accelerating against 39% decelerating, and that number was an artifact of our own design. A gap cannot be shorter than one week, so for a triangle whose median gap is two, “accelerating” was arithmetically unreachable; 61 of 80 triangles could not report it at all and every one returned zero. Any statistic that cannot report one of its answers is not a test. The tool now prints *reachability* — how many units could have returned each verdict — before it prints any verdict, and we recommend that practice to anyone building something similar.

10.2 The census: what the lattice is, before asking what it does

Of 441 triangles, 425 carry era weeks, and **225 of them (53%) exceed three times their own coin-flip benchmark** — the level a triangle of the same activity but random signs would reach. Over half this lattice genuinely winds rather than flickering.

Two structural facts came out of the same run, and both change how everything else should be read.

Fifteen populated chart labels exist (\$1) and only six ever wholly contain a triangle. Of 425 loops, 169 (40%) sit inside a single block and 256 (60%) straddle two or three. The smaller blocks are too small for the lattice to close a loop inside them. The majority of this instrument lives on boundaries rather than within them, so the hand-off loops between blocks are not a minority curiosity — they are the bulk of the network.

Activity and one-sidedness are completely independent. Splitting the triangles at the medians gives 27%, 24%, 23% and 27% across the four quadrants, and the correlation between charged weeks and drift is **−0.016**. How busy a triangle is tells you nothing about whether it goes anywhere. That kills one clean hypothesis outright: there are no separate populations of continuously-releasing and steadily-storing patches visible in the time-averaged record. A patch that spent a decade doing one and then switched would have both regimes averaged into a single middling number, so the regime-switch version of that idea survives untested.

Every triangle: how big it is against how hard it winds

CASCADIA_CENSUS v1.2 · run 20260818_211537 · census, no test · x is log area, y is |net|/charged-weeks · dashed = what a random-sign churner gives




Figure 11. Every triangle: how big it is against how hard it winds, with the random-churner benchmark drawn beside it. Drift decreases with area — bigger loops are churnier. (CASCADIA_CENSUS v1.2, run 20260818_211537.)

The 25 largest triangles — the long reaches

CASCADIA_CENSUS v1.2 · run 20260818_211537 · these are the loops wide enough to enclose offshore ground; every short inland triangle wraps crust that is not the fault




Figure 12. The twenty-five largest triangles — the only loops on this lattice wide enough to enclose offshore ground. (Same run.)

10.3 Blocks do not turn as bodies

If the winding were a block moving as a body with a wandering rotation rate, every triangle inside it would ride the same rotation and would have to agree week by week. If instead the winding is a local mismatch — a loop failing to close under a deformation the ground cannot satisfy — neighbours need not agree at all.

Measuring, per block per week, how lopsided the charged triangles are against a fair-coin resample at the same count: **every group came back negative**. Individual blocks at -0.19 , -0.13 , -0.33 and -0.32 ; seams at -0.17 through -0.30 ; the whole margin at -0.05 . Nothing anywhere near positive. Blocks are not turning as bodies.

The direction of that comparison is the finding; its magnitude is not physics, and we say so. On an oriented mesh two triangles sharing an edge traverse that edge in opposite directions, so a large phase step on a shared edge pushes one face positive and its neighbour negative. Adjacent faces anti-correlate *by construction* — the same

cancellation that makes discrete Stokes work — and a fair-coin null knows nothing about shared edges. A null able to price the magnitude would have to permute the phase data and rebuild the charges, preserving edge structure. We have not done that.




Figure 13. Coherence by block against the fair-coin null. Read the sign, not the size. (COHERENCE_TEST v1.1, run 20260818_224138.)

10.4 What does winding mean? The question we could not close

Over half the lattice winds more than chance, and we still cannot say what winding *is*. A patch creeping steadily winds steadily and stores nothing; a patch locking up hard also winds steadily. The measurement is agnostic about its own meaning, and no amount of internal consistency resolves that.

The discriminator requires a second instrument, and not one derived from the same GPS — which rules out published locking models, since those are inverted from onshore geodesy. Tremor epicentres come from seismometers: different instrument, different physics, genuinely independent. And tremor is the field's most direct marker of ground that is actually moving, so the sign reads directly. That gives us two competing readings of what a winding is, and they are worth naming because they make opposite predictions. Call the first **lubrication**: the winding sits where tremor is dense, tracking ground that demonstrably slides, so winding is the signature of a fault surface letting go. Call the second **loading**: the winding sits where tremor is absent, on ground that is locked, so winding is strain piling up against a patch that will not move. Lubrication predicts winding and tremor density to rise together; loading predicts them to run opposite. One correlation separates them.

We pulled **211,581 tremor detections** across 25 windows, one per clean onset, from the public catalogue. 62.7% fall inside at least one triangle and 219 of 315 usable triangles contain detections — better overlap than we feared for a land network reading a source 30 to 45 kilometres down.

The rank correlation between a triangle's drift and its tremor density is +0.056, against a null band of −0.111 to +0.109 and a permutation p of 0.32. Splitting at the February 2018 detector change gives +0.068 before and +0.041 after, so the answer is not an artifact of the catalogue changing under us. The same correlation against *activity* rather than drift is −0.081, also nothing.

Both readings are unsupported. Lubrication predicted a positive correlation; loading predicted a negative one. Neither appeared. The discriminator ran, and it discriminated against both candidates. Where a triangle winds one-sidedly has nothing to do with where the ground slips.

Two things make that null stronger rather than weaker. Density is detections per unit area, so a small triangle wholly inside the tremor band scores high while a large one straddling its edge is diluted — and the census measured that smaller triangles *also* drift higher, 0.500 in the smallest size band against 0.394 in the largest. Both effects push

toward a spurious positive, and it still did not clear the null. And the tool refuses to report “no relationship” when either variable lacks spread, because a rank correlation with no variance returns zero and looks exactly like a clean null; both varied here, across 237 distinct drift values and 220 distinct densities.




Figure 14. Drift against tremor density, one dot per triangle. Rising to the right would mean the winding tracks slipping ground; falling would mean locked ground. It does neither. The vertical stack at the left edge is the triangles containing no detections at all, drawn at the 0.0001 floor a logarithmic axis requires — zero cannot be placed on a log scale. (TREMOR_DRIFT v1.1, run 20260819_060144.)

The caveat that could still hide a relationship. Tremor is deep and these are surface loops. The surface expression of a deep slip patch is smeared over tens of kilometres and offset updip, so pairing a triangle with the epicentres strictly inside it is a sharp instrument aimed at a blurry source. Counting detections within a radius of each triangle’s centroid would mimic the broad surface expression without touching a GPS-derived slip model, and stays inside the independence fence. We recommend it as the first thing to try.

And a third possibility now has to be on the table. If the winding neither tracks slip nor avoids it, some of it may not be reading fault-zone processes at all — hydrological loading, reference-frame residual, or station-local effects surviving detrending and common-mode subtraction. This paper already treats the subtracted common mode as a nuisance we learned to measure. Fifty-three percent of triangles beat their own noise floor, and that winding is not co-located with slip. Those two facts sit together uncomfortably and we do not resolve them.

10.5 What the network’s shape decides, before any data is collected

The most transferable thing this work produced is not a result. It is that **a network’s geometry determines which questions can be asked of it, and no amount of data quality substitutes.**

We designed an analysis that blinds a seam — removing every loop touching the boundary between two blocks, then reading the seam’s behaviour from a closed ring of witnesses that never touched it, recovering the blinded interior by Stokes from stations that could not have been coached. It is a clean design and Cascadia cannot support it. Across 19 seams and 25 onsets, **not one produced a walkable rim.** A seam runs *across* a coastal strip, so blinding it cuts the strip in two and reaches the edge of the map at both ends. The boundary broke into two pieces, or pinched through a single station — the same station, P736, on every southern event for sixteen years — or came back as a clean ring that nonetheless carried blinded stations because the blinded patch had no lattice outside it.

On a coastal strip network a seam cannot be encircled. That is a property of where the network was built, not of the data or the method, and it will be true for any margin shaped like a ribbon.

But the measurement survives intact one scale down, and there the network is far more capable than the seam failure suggested. Every station sits inside a fan of triangles whose other corners form its *link* — a ring of neighbours. Where that fan closes, the station can be deleted from the data entirely and its own fan recovered from

the ring alone, because the spokes from centre to rim cancel in pairs and the winding walked around the link equals the sum of the faces inside it.

211 of 228 stations have a closed fan. Only seventeen can never be interrogated from outside, whatever the data does, and each is named with its reason: QUAD has a fan of two triangles, P157, P327, P339, P446, RMDB and TPW2 have three, and BCWS, CHLW, CME5, GLDR, LINH, P148, P164, P389, QMAR and WAYA have fans that do not close, each opening at a named pair of neighbours. CME5 opens at GLDR and GLDR opens at CME5 — the Cape Mendocino pair, where the lattice runs out of ground. Ring sizes run from four to nine neighbours, median six.

The theorem was then checked on the real data rather than asserted. Across 17,950 station-weeks carrying both a walkable ring and a fully defined fan, the ring walk equalled the face sum in 17,950 of them. Zero disagreements. What binds the measurement is not the geometry but the amplitude gate: a further 70,767 station-weeks had a perfectly good ring with at least one corner below the gate, so roughly four fifths of the available record is unreadable for reasons that have nothing to do with the loops.




Figure 15. Which stations can be interrogated from outside. Blue: the fan closes, so the station can be blinded and its own fan recovered from its neighbours. Red: the fan is open — coast, inland edge, or a gap — and no ring exists, ever. (STATION_RING v2.3, run 20260819_174258.)

And with that instrument we went looking for tattlers — stations that give a coming onset away by winding differently in the weeks before it — and there are none. Each eligible station was blinded in turn and its recovered fan winding read over the twelve weeks before each onset, scored on the walk alone so the blinded station's own data can never leak back in, ternary only, banded by that station's own shuffled-clock null. Across 52 scored stations on the clean set the fraction of onsets reading non-background has a median of **0.26** against roughly 0.32 expected by construction; across 73 stations on the full set including every contaminated event, **0.25**.

The shape of that null is the shape that fools people. Every station at or above 0.60 has between three and six scored events: BCSF and P442 both sit at 0.75 on four events apiece. Three of four non-background occurs by chance about ten percent of the time per station, so across 52 stations one expects roughly five at that level from noise alone — and nine appear. Meanwhile the well-sampled stations run from 0.00 to 0.50 with a median of 0.267, and on the full set the whole range collapses from 0.75 to 0.46. **The better a station is sampled, the closer it sits to background.** A real effect converges as sampling improves; this one dissolves. BCSF is not a witness that knows something, it is a station that was asked four times.




Figure 16. The ring-walk lag profile — when each ring of neighbours lights relative to onset, with the shuffled-clock null beside it. (BLIND_RING v2.3, run 20260819_174252.)

Two bounds belong with that null. We asked one question — does the recovered fan winding depart from background in the twelve weeks before an onset — and a station could carry information in a form we never asked for. And with three to fourteen scored events per station, and four fifths of the record gated away, we could only have caught a station that spoke on *most* occasions. One that whispered on a third of them would be invisible here.

10.6 The waterfall, the drops, and the loose governor

One more post-hoc observation, offered as a picture and a first measurement. The rhythm measurement handed us an irreducible pair of numbers: once a triangle completes a winding, it winds again *the same way* after a median of two weeks, but a *reversal* — the sign flipping — waits a median of eighteen weeks (in the training era alone, one and seventeen). The two-week figure is a floor, not a value: positions are collapsed to weekly medians before anything else happens, so any faster return is invisible at this cadence, permanently. Two clocks, one fast and one slow, and no metronome anywhere: not one of the 371 triangles the rhythm census could score ticks regularly, and ninety-one percent are bursty. Read as a point process — each completed winding an integer tick, a timestamp and nothing more — the lattice behaves less like a clock and more like an old engine with a heavy flywheel and a loose governor: repeated strokes in one direction on the fast beat, momentum carried through the silences, surging and loping rather than idling smooth, with the rare counter-stroke on the slow clock.

We then asked the engine question directly: does the three-state cycle *circulate* — the signature of energy flowing through rather than noise idling? Counting every transition between adjacent readable weeks — 84,904 of them — the cycle circulation measured **$C = -14.0$ against a shuffled-sequence null band of -67.0 to $+63.3$, two-sided $p = 0.6965$** (WATERFALL_MAP v1.1, run 20260821_200450): squarely inside the band. **The cycle does not circulate.** Detailed balance holds, the state walk is time-reversible, and the thermodynamic-engine reading of the flywheel picture is falsified in this channel; even per triangle the largest circulation anywhere is 4.67, noise-scale, and neither the dipole nor the big winders appear in the top ten. The transition record still carries structure, and two pieces of it are worth printing. First, the flywheel is real but has no preferred direction: a charged triangle continues its sign into the next readable week with probability 0.639 (+) and 0.641 (–) — momentum without direction — and quiet stays quiet at 0.932. Second, **the board has no reverse gear**: of 13,569 departures from a charged state, only 129 —

under one percent — flip sign directly. A reversal routes through zero essentially always; the fall goes through the crest. The ternary zero as the crossover state is not a labelling convention here — it is how the transition record actually moves. And the same census closes a loop with Part II's matcher-free curve: the network winds in shared episodes months long — T5b's elevated response curve is that co-clustering measured directly — which is exactly what episodic tremor and slip should look like to a counting instrument.

Two cautions keep the picture honest. The record's zeros mix two different silences — a loop truly reading zero circulation, and a loop muted by the amplitude gate — and roughly seven in ten triangle-weeks are the second kind: 68.9% gated, 25.9% defined zero, 5.3% charged. And the rhythm carries no warning as measured: it did not change before onsets.

A conservation-shaped question remained. If one deep source feeds many relief valves, the *count* an episode dispenses might be steadier than its loudness. This instrument discards magnitudes by design, so a true energy ledger is impossible here — but a count budget is not. For every catalogued tremor episode we counted the network's completed windings in the seven weeks from onset (WINDING_BUDGET v1.1, run 20260821_190021), per segment, against a null that rotates the segment's whole episode clock rigidly — every spacing and cluster preserved, only the alignment destroyed — a null that is conservative by construction on a quasi-periodic clock, because some rotation offsets partially re-align with the real seasonality. **The south delivered: SOUTH (Klamath), the matched bench with 257 usable repeats, reads CONSISTENT DISPENSER — episode-to-episode variability (CV 0.4201) below its own rotation-null 5th percentile (0.4236).** Six fences go to print with it. It is thin — below the band by 0.0035, a boundary crossing. Three segments were tested and one crossed at the 5% level, roughly a fourteen-percent chance of one crossing by luck alone — suggestive, not confirmed — though the crossing landing in the segment with by far the most repeats (257, against 137 in the north and 40 in the central) is what a real effect would do. The north leans steady but earns no verdict (CV 0.4316 against its own 5th percentile of 0.4157), and the central segment, with only forty usable episodes, sits inside its band. The budgets are *not elevated*: episodes dispense a mean of 160-164 windings, about what any arbitrary seven weeks dispenses (~164 at the record's 5.3% charged rate) — the finding is consistency, never quantity. Dose-response is null in all three segments (Spearman +0.10, -0.03, +0.08 against null bands reaching 0.11-0.15): bigger deep events do not dispense more windings, rhyming with the tremor-drift spatial null and this instrument's magnitude-blindness by design. And the home-share medians — the north keeping 0.443 of each budget in its home region against its 0.374 share of the lattice, the central and south at 0.248 and 0.276 — are a post-hoc observation only.

Whether this two-clock engine is how a subduction margin maintains its steady state — energy nudged through the network drop by drop, released and reorganized on the slow stroke — is speculation we flag as such. Two instruments we did not build are open invitations. The fuller count-budget study — per-valve budgets, conservation across neighboring episodes, the ledger read on other margins' episode catalogues — is a weekend's work on top of the deposited tools. And the honest next instrument would strip the map coordinates entirely and let the triangles seat themselves by shared behavior — a co-behavior embedding in which “near” means “moves together,” whatever the kilometres say. Let the raindrops fall where they may.

10.7 What we are asking of anyone reading this

Not agreement. Three specific things.

Check our work. Every tool that produced these numbers recomputes the canonical record from source and halts on a single mismatch, every null carries the same freedoms as the measurement it prices, and every design was stated before the data was touched. The tools refuse to write a report and a data file that disagree with each other. All of it is deposited. If a number here is wrong we would rather learn it from you than from history.

Grab a theory and test it. We have three findings we cannot explain: the nine-to-one asymmetry between winding onward and winding back, a margin with no metronomes anywhere on it, and a winding field that has nothing to do with where the ground slips. We offer guesses above and label them as guesses. If one of them is yours to test, take it.

Run it on your own margin. The instrument is not ours and the answer here is one margin's answer. A network with area rather than length will reach questions Cascadia refused us; a strike-slip margin will need a shear channel before the result means anything; a dense catalogue will settle in a season what took us a decade of record. And a flat answer from your network is as valuable as ours was, because it is the second one — and two independent flats are worth far more than one.

Interlude — Skiball, and the Game the Data Asked to Play Instead

We went to the geodesy Midway to play with our data, and everyone was lined up at the skiball machine. That's usually our line too. Skiball is the honest name for how a precursor campaign works: you build a hypothesis as an equation of proposed coefficients and variables, you take your throws, and you chase the 100 ring — a signal that leads an event. We did exactly that. And our fiber-bundle network map refused to play by the old rules. It was not silent — it was very forthcoming, a regular chatterbox of data — but through the eyes of a skiball player, it failed. Every ball question came back empty. So we said, huh, that's weird; let's report it, and maybe someday someone will explain it a little further for the cheap seats like ours in the back of the theater.

Then, rereading this draft for the hundredth time, we hit our own opening line: *data needs to speak, and it speaks through visualization and metaphor — analogy is not decoration here, it is the method's front door*. We had just spent a section humbly explaining that forcing our geometric projections onto the data is what killed the original campaign. The data was still trying to teach us. So we asked it to name a better game, and the game it named is pachinko. A fair pachinko board is perfectly random: every ball's path is chaos, every pocket equally likely. But parlors bend the pins, and a bent board does something strange — each individual ball stays unpredictable, while the *pile* becomes stable, repeatable, and set entirely by the pins. Physicists call the bent pins quenched disorder. The crust has been bending its own pins for billions of years — sutures, dead faults, compositional seams, asperities — which may be exactly why earthquakes keep happening at the same addresses. We were playing the wrong game to model this data.

One correction before the frame does any work: our board runs upside down. A pachinko ball falls, and the thing we are watching rises. So hold the picture the way it arrived — a tall glass packed with ice, fizzy water poured over it, and the fizz climbing. Something happens deep beneath the margin. A bubble does not choose its route; it takes whatever path the ice leaves open, and it surfaces where the cubes decide. The ice is the crust's inherited structure, the bubbles are what the deep event sends upward, and the stations and the triangles they form are the slots at the top where the count lands. Any single bubble is hopeless to predict. The pile is not. That is the whole pachinko claim, upended and poured into a glass, and everything Part II tests is one or another version of it.

We should be careful about what we call the thing at the bottom. Conventional wisdom names these episodes tremor and slip, and the catalogues are built on that reading. What this instrument can confirm is narrower: a deep event, its timing taken from a public catalogue of tremor episodes, propagating upward through structure — and an instrument that reads the structure rather than the event. We never measure slip. We never measure the size of anything, because this channel is magnitude-blind by design, which is why §10.6's null dose-response is a measurement rather than a gap. We take the dates and decline the mechanism, for the same reason §10.4 declines to say what a winding is. Informally, and with apologies to the seismologists, we have come to call the deep events “burps.” The word is honest about how little we know: something down there let go, and we are reading the fizz.

Read our whole scoreboard through the pachinko frame and it stops being mixed and becomes one sentence: **the ball channel is empty and the pin channel is full**. Every trajectory question — does the ordering lead, does the rhythm accelerate, do the tattlers speak before onsets — came back a clean null. Every board question came back positive: the standing texture held its spatial configuration blind through the vault, the dipole kept its address for eras and survived by name, the rogues live at fixed seats, the southern episodes replay on the same seam over roughly sixteen repeats. The instrument we built turns out to read board geometry, not ball flight — a pin map, not a fortune teller. Two honest fences: the parlor bends its pins on purpose, and we make no such claim for the crust — inherited structure, not intent; and “wrong game” means wrong for *this instrument on this margin*, not that prediction is impossible for every instrument on every channel. But if you came here hoping we'd guarantee the 100 ring, we can't, and we now suspect the machine was never going to pay out that way for anyone. What we can hand you is the map of the bent pins.

Part II — The Ascent: the Pachinko Board Made Falsifiable

Part I descended: a hypothesis shaped like an equation, chasing a precursor on an event, down to a flat answer honestly earned. §10 walked the strange ground we found on the way down. Part II climbs back out. If the crust is a pachinko board whose pins were bent by billions of years of inheritance — stable addresses shaping outcomes while every individual trajectory stays unpredictable — then the claim splits into two separable channels, and each can be tested. Five tests were frozen in PACHINKO_PROTOCOL v0-3 (sha 1a5286f4de6f), ratified as written before the suite tool existed. The qualitative observations that motivated them — the lateral streaks, the negative space, the clustering — were made post hoc, by eye; the tests were specified before any of these numbers existed. Eyes had touched the picture; no eyes had touched the numbers. One run then exposed one of our own instruments as broken, and Amendment A5 (v0-4, sha fb8c33ca8188) repaired it with honest dating: designed after the spent-record run, before any sealed-era data was pulled or touched — the vault stayed shut throughout, so the repaired instruments are legitimately frozen for the registration that closes this paper. The run of record is PACHINKO_SUITE v1.2 (sha a426615fe171), run 20260821_191042, seed 20260821, integrity gate 441 of 441 exact.

The two families are reported separately and are never merged into one score. THE BOARD (T3, T4) asks where stable behavior lives: its statistics are decade-scale averages per address, timing-blind by construction, so no arrival smear, lag, or burst tail can enter them. THE FLOW (T1, T2, T5) asks how energy moves: its statistics are timing-only — same-week agreement, lag panels, lead-lag stability — and never contribute to any address's behavior card. A reader may believe the board and doubt the flow, or the reverse; nothing below combines them.

The registration: what was frozen, and when

The process is the vault pattern this paper already proved once: ratify the protocol as a written document, build and pin the suite tool with selftests and the write contract, one double-click on the operator's machine, and the results ship in the paper as-is — works, partially works, doesn't work, or null. The master seed family is 20260821; every random draw derives from it and is printed in the run report. Two frame fences were stated once in the protocol and inherited by every test. THE WEEKLY SMEAR: positions collapse to weekly medians before anything else happens, so any delay shorter than one week is invisible at this cadence, permanently — "same week" means "within the same 0-7 day bin," and the fast clock's two-week median is a floor, not a value. MIGRATION vs COMMON CAUSE: a lag between two stations can mean energy genuinely traveling a channel or two valves answering one deep cause on different local response times, and this instrument cannot tell them apart from a single event — the discriminator is repetition, so stability of a lead-lag across many bursts, not its existence, is the only claim shape allowed.

What was frozen, in one breath each. T1, coherence: when several loops speak in the same week, do they agree on direction more than timing-scrambled ground would — priced against the expensive, honest null, a per-station circular time shift that rotates each station's weekly residual series by an independent random offset and recomputes charges through the exact canonical recipe, preserving every station's own statistics and the lattice geometry while destroying only cross-station timing. T2, the audibility control: a noisy week un-mutes many loops at once, so agreement must be re-read within audibility quartiles before it counts. T3, the negative-space census: each triangle's quiet-plateau lengths against the memoryless distribution implied by its own one-step stay-quiet probability, simulated on its own gating mask. T4, the pin test and the headline: four-number behavior cards per half-record, split at the vault wall, the studied era's map required to predict the era the design never touched, against shuffled addresses — with the direction pre-declared: the pachinko frame predicts positive correlation, and failure to find it counts against the frame and is reported exactly as loudly. T5, lead-lag consistency: do pairs of valves answer the same deep burp with the same delay, burst after burst. All five were specified and frozen in v0-3 (sha 1a5286f4de6f) before the suite tool existed.

The honest dating is part of the registration. The qualitative observations that motivated these tests were post hoc, by eye. The five tests were frozen before any of their numbers existed. The first canonical run then exposed T5-v1's null as broken, and Amendment A5 (v0-4, sha fb8c33ca8188) added its two replacements — T5-v2's matched-zeal null and T5b's matcher-free response curve — designed after the spent-record run and before any sealed-era data had been pulled or touched. The sealed vault stayed shut through all of it, which is what makes the repaired instruments legitimately frozen for the registration that closes this Part. The full protocol ships in the deposit crate as the written document the tools were built against.

The board: the pins hold

T4, the pin test, is the headline. Each triangle gets a four-number behavior card per half-record: drift (how one-sided its winding is — absolute net over charged weeks), activity (how often it is charged at all), persistence (the probability it winds the same way next week), and its median quiet plateau. The record splits at the vault wall — half A is wk500-1094, the era we studied; half B is wk1095-1383, the era our design never touched — so the pin map drawn on the studied era must predict the era we were blind to. Across 351 eligible addresses the combined rank correlation between the two half-maps is **0.2697 against a shuffled-address 95th percentile of 0.0940 — the pins hold** — and all four features hold individually: activity 0.5393 (null 95th 0.0926), persistence 0.1802 (0.0924), drift 0.1494 (0.0822), plateau 0.1032 (0.0897). The pre-declared sensitivity line at the equal split (wk941/942) reads 0.2499 against 0.0837 — the verdict does not depend on where the knife fell. The straddle census puts the seam's contamination on the page rather than in an argument: 16 charged bursts and 94 zero plateaus cross the wall. The interpretation fence travels with the result: two halves of one margin share the same slow history, so "the addresses hold" means persistent addresses on this margin over this record — the instrument cannot distinguish a pin bent into the rock from a slow process that keeps expressing at fixed addresses. Both are the board, not the ball.

T3 asks whether even the quiet has addresses. Per triangle, the lengths of its zero plateaus are compared to the memoryless distribution implied by that triangle's own one-step stay-quiet probability. The center of the quiet is memoryless — the observed and simulated median-of-medians are both exactly 1.0 — but **29 of 357 evaluable triangles are longer-tailed than their own memoryless simulation against a chance expectation of about 17.9, and zero are shorter-tailed.** The typical silence has no memory; a minority of addresses own long silences. The wide quiet swathes in the raster live at seats, which is itself a pin statement.

The flow: a whisper at lag zero, and nothing that travels

T1 measures whether loops that speak in the same week agree on direction. Over 869 qualifying weeks (at least six charged triangles; floors of four and ten are printed as sensitivity, never verdict), the median sign imbalance is **0.0833 against a circular-shift null band of 0.0612-0.0769 — outside the band, thin, and real**; 10.1% of qualifying weeks exceed their own null 95th percentile against the 5% chance rate. The pre-registered lag panel answers where the synchrony lives: the statistic falls immediately off lag zero (0.083 at $k=0$, then 0.059-0.064 across $k=1..8$), and the multiplicity-priced max-over-lags verdict is driven entirely by $k=0$. Whatever synchrony exists is same-week only.

T2 is T1's control, and T1 survives it. A noisy week un-mutes many loops at once, so “everyone talking together” can secretly mean “everyone audible together.” The weekly charged count does track the audible count (Spearman 0.8732, inside its own null band of 0.7411-0.8782 — gating co-varies exactly as timing-scrambled ground would), but the decisive line is coherence by audibility quartile: **0.143 in the quietest weeks falling to 0.053 in the loudest.** The agreement is strongest precisely where the fewest loops are audible — the opposite direction from the one the confound would push. The whisper is not a gating artifact.

T5-v1 stands here as a documented instrument failure, claiming nothing. Its greedy matcher scored against a shuffled-pairing null returned 35,390 of 38,814 pairs “exceeding” chance of roughly 1,941 — a number that carries no information, because the null was broad by construction while the matcher concentrates lags for any bursty pair. The run's own exemplars carry the smoking gun: modal lags sit only at zero (14 of 25) or at the window edge of 7-8 weeks (11 of 25), never between, and hub triangles recur (CABL-GLDR-CME5 appears in 9 of the 25) — the signature of the search, not the ground. This is the failure Amendment A5 repaired, and it is kept in print because a built-instrument lesson is part of the record.

T5-v2 re-asks the question with a null that beeps on grass at the matcher's own rate. Same pairs, same matcher, same statistic; the null circularly time-shifts one member's entire arrival series — preserving its inter-arrival structure and burstiness exactly — and runs the identical matcher on the shifted pair. The answer: **1,769 of 35,494 evaluable pairs exceed their own null 95th percentile — 4.98% against the 5.00% chance expects. Dead on chance.** 3,320 pairs were unevaluable and are counted separately; the dependence fence holds (pairs share triangles, so the count is read qualitatively, never as a binomial p-value); and the geographic breakout is flat — within-south 5.4%, within-north 4.6%, cross-divide 5.1%. The 91% collapsed to exactly nothing, which means two things at once: the repair worked, and there is no detectable stable pair-delay channel at this instrument's resolution. The migration question is answered no at weekly cadence — with the weekly-smear fence stated plainly: any delay shorter than one week is invisible here, permanently.

T5b removes the matcher entirely. Ring every bell once and listen to when the rest of the network rings around it: for every completed winding, the other triangles' windings are counted at each lag from -8 to +8 weeks (the source and its shared-edge neighbors excluded, since their anticorrelation is mechanical), averaged per event over 20,511 events. The registered max-over-lags verdict reads “exceeds” — per-event mean 27.0457 at lag 0 against a null 95th of 23.0669 — but the honest reading is in the curve's shape. The entire curve sits about 15% above its null band at every lag (26.5 at the ± 8 edges against a band top near 22.9-23.1): broad co-clustering, the network winding in shared episodes months long — which is what episodic tremor and slip looks like to this instrument, and which a circular shift destroys by construction. Riding on top is a small, genuine lag-0 bump — 27.05 against 26.74 at its immediate neighbors, about 1.2% — an independent confirmation of T1's same-week whisper by a statistic that shares none of its machinery. One fence is structural and the paper must say it: this curve is symmetric by construction, $R(k) = R(-k)$ exactly, because every event pair is counted in both directions. It cannot distinguish lead from lag. It is the network's pair-correlation function, and it is read as one.

The two verdicts, kept apart

The board channel is loud: the pin map drawn on the studied era predicts the blinded era on all four features, and even the silences have addresses. The flow channel whispers: a thin, real same-week agreement (T1, confirmed independently by T5b's bump), broad episode co-clustering (T5b's elevation), no stable pair-delay channels anywhere (T5-v2), and no audibility artifact underneath any of it (T2). Every individual trajectory remains unpredictable; the addresses hold. That is the pachinko frame's own prediction, met on the channel it named and absent on the channel it never promised — and the two statements are left standing separately, as registered.

What the Whole Ride Found

Two arcs, one instrument, one scoreboard — and the scoreboard only reads straight when both halves are laid side by side. Everything below is already in this paper, with its run stamp attached where it was earned. Nothing here is new. This section exists because a reader should not have to reassemble thirty pages to find out what happened.

The descent. We came to Cascadia with the ordinary hope: build a lattice, read its windings, find something that leads an earthquake. The lattice built cleanly and the windings were real — 84,904 of them across the studied era. The lead never came. The precursor campaign ended flat, and it ended flat honestly, with its nulls priced before its

questions were asked and its design frozen before its numbers existed. On the way down the instrument did show us things worth keeping: a silent traveling front caught twice in the training record, at wk1146 and again at wk1151, migrating north at two to five kilometres a day — slow-slip speed, read from GPS geometry alone, against a contrast specimen at wk1360 that lit 91 stations in a single week with no buildup and no migration. Two morphologies, one instrument that tells them apart. A witness pool of 28 stations elected before the vault and frozen by name. A dipole, on one shared edge in the Olympic coastal group: one triangle standing pure +120 of 120 charged weeks against its neighbour's pure -111 of 111, found by the first instrument that could have seen such a thing. Four instruments queued up to convict it of being one bad station shared between two triangles, and none of them could — and that failure measured something better than a conviction would have, because the residual field turned out to carry standing spatial texture, map-wide and persistent, with the dipole sitting in it as a knot where the frozen texture happens to wind. And that standing texture held its spatial configuration blind through the vault opening — the one result that survived the descent intact.

The low ground. The addendum is where the strange findings live, and they are strange in a consistent direction. This margin has no metronomes: not one of the 371 triangles the rhythm census could score ticks regularly, and ninety-one percent are bursty. Blocks do not turn as bodies — every group came back negative against a fair-coin null, so the winding is a local mismatch rather than a rotation carried by a block. The winding field has close to nothing to do with where the ground slips. Circulation does not circulate: the network's net signed transport measured -14.0 against a null band running -67.0 to +63.3, $p = 0.6965$. And the record has no reverse gear — of 13,569 scored transitions only 129 cross directly from winding one way to winding the other, under one percent, with reversal routing through zero instead. The count budget added the one conservation-shaped result we have: the southern segment dispenses windings at a steadier rate than its own rotated clock predicts, CV 0.4201 against a null of 0.4236, which is consistency and explicitly not quantity.

The climb. The Interlude changed the game and Part II froze the new one before scoring it. The board channel came back loud: the behaviour map drawn on the studied era predicts the blinded era at $r = 0.2697$ against a null ceiling of 0.0940, on all four features at once, and at 0.2499 under an equal-split sensitivity that removes the split's own freedom. Even the silences have addresses — 29 of 357 stations own significantly long quiet runs against 17.9 expected by chance, while the centre of the quiet distribution stays memoryless. The flow channel whispers and no more: a thin but real same-week agreement at 0.0833 against a null reaching 0.0769, independently confirmed by a pair-correlation bump that shares none of its machinery; broad episode co-clustering; no stable pair-delay anywhere, with the matched-zeal test landing at 4.98% — dead on chance; and no audibility artefact underneath any of it. One of our own instruments broke in the process and is printed as broken, exemplars and all, because a documented failure is worth more than a quiet deletion.

The teaching items, and what each one is worth

1. **A registered null is a result.** The precursor campaign found nothing and that nothing is the paper's most reusable finding — because it was priced against a null built before the question, on a margin whose properties were published alongside it. A second flat answer from another network is worth more than this one was, and that is the point of publishing the first.

- **Freeze the design before the numbers, then print the failures exactly as loudly.** T4's direction was pre-declared in writing, with the clause that failure would count against the frame. The flow channel then came back empty and was printed empty. A frame that can only win is not a frame.

- **Price every claim against a null with the same freedoms it had.** The question was never “does this margin cluster” — everyone knows it does. It was whether the structure exceeds what the same record produces when its addresses are shuffled. That is a question with an available no.

- **A one-shot vault beats any amount of re-analysis.** The sealed era was sealed when the design was frozen, so nothing could be tuned toward it. Two of the first vault's predictions validated blind. No amount of careful post-hoc work buys what one honest blind test buys.

- **Check the artifact, not the intention.** Every tool here recomputes the record from source and halts on a single mismatch. That habit caught a miscounted witness pool, a mis-scoped denominator and a stale block of assembly notes inside this build alone — each one found by asking a file rather than a memory.

- **Declared blindness is a feature.** This channel discards magnitudes on purpose, so the null dose-response is a measurement rather than a shortcoming. An instrument that says in advance what it cannot see can be trusted about what it can.

- **When the data refuses your game, change games — and then freeze the new one before you score it.** The pachinko reframe arrived post hoc, by eye. It was written into a protocol and ratified before the suite tool existed, so the reframe could be tested rather than merely admired.

Where the measurement stops and the guess starts

The honest boundary matters more here than anywhere else in the paper, so we will draw it explicitly and against our own favourite idea.

What was measured is that spatial behaviour persists across a wall the design never crossed, at a stated rate, against a null carrying the same freedoms. What is narrated is that the crust bent its own pins across geological time and the lattice is reading them. Those are not the same claim. Site quality, monument behaviour, local noise environments and network geometry all produce spatially persistent behaviour too; the audibility control and the grouping sensitivity close some of that door and not all of it. The pachinko frame is a frame — its job in this paper was to generate falsifiable tests, and it earns its keep by having generated five of them across two channels, one whole channel of which came back empty and was reported empty.

We name the obvious objection ourselves: that a frame invented after looking at pictures, and then confirmed, is a target painted around the bullet holes. The defences are on the record rather than in this paragraph. The tests were written before their numbers existed. The direction that would count against the frame was declared in advance. Half the frame's predictions failed and were printed. The strongest single result was scored against data that was sealed when the design was frozen. And the campaign that preceded all of it died flat and was published dead. A painted target produces no misses; this paper is most of the way made of them.

What stayed open

We could not close what a winding *means*. Section 10.4 says so at length and we have not improved on it since. We do not know whether the inland micro- blocks wind as one body or merely appear to; the coherence test says blocks in general do not, and the direction of that test is the finding while its magnitude is not physics. The silent front is a specimen family, not a statistic — $n = 1$, arguably 2 — and it needs a second margin or a longer record before it means more than method-plus-example. The count budget wants the same ledger read on other margins' episode catalogues. And the raindrops question — whether stations that behave alike sit near each other once you stop assuming geography decides behaviour — is a weekend's work for whoever wants it. Every one of these is an invitation, and every one of them is cheaper to answer than it was to ask.

Vault II — The Sealed-Era Registration, Scored by Anyone

The cache froze at wk1383, in early July 2026. Everything after it is sealed by construction — never touched by our eyes, our tools, or our choices — and the sealed window grows a week wider every week, for free. This section is the paper's loaded experiment, handed to the reader: five predictions, frozen now, scoreable by anyone with the deposited crate and a fresh data pull, and binding on us in whatever future scoring we publish. One naming note so the ledgers stay distinct: V2-1 through V2-5 are the second vault's predictions, numbered as the protocol froze them; Part I's event-contrast claim V2 is a different object and re-enters through §9's front door.

Frozen by name. Every structure below is scored as the printed list it already is: the 28-station witness pool by name (the frozen protocol's text says 27 in one mention — a miscount of its own printed list, and the deposited list of 28 names governs, which is exactly what frozen-by-name means), the dipole as its two named triangles, the fifteen block charts as the deposited clusters.json membership, the lattice as the deposited triangle file. The sealed era supplies new motion for the same old names and is never allowed to redraw the structures — a scorer who re-derives the clustering or the lattice on fresh data is grading a different exam. All defining files ship in the crate sha-pinned, and the scoring recipe names them.

Fixed scoring windows. H1 scores wk1384-1435 exactly; H2 scores wk1384-1539 exactly. Open-ended scoring invites the quiet cheat of peeking repeatedly and publishing the favorable pull; the fixed window is the optional-stopping fence, and it binds us as much as anyone.

Roster attrition. Stations die. Scoring runs on the surviving intersection of the frozen roster and the fresh pull, with the survivor count printed. A prediction cannot be dodged by roster shrinkage, and it cannot be failed by it either.

Horizon H1, wk1384-1435 — one sealed year, the earliest date each promise has teeth: roughly 1,030 fast-clock gaps, roughly 151 reversals, block charts averaging about 30 triangles.

V2-1 · The two clocks hold. The sealed-era same-way median falls within 1-3 weeks and the reversal median within 12-24 weeks. The bands are measurement-backed, bootstrapped from the engine census run's own gap lists (run 20260821_070732) — one sealed year's sampling ranges are same-way 1-2 and reversal 14-22, plus margin — and the miss semantics are printed in advance: a same-way median of 4 or more means the fast clock's rate halved; a reversal median outside 12-24 means the slow clock moved by more than a third. Either is regime change, not noise.

V2-2 · Pool retention repeats. The frozen 28-station witness pool remains over-represented among the sealed era's most active vertices, above the shuffled-roster 95th percentile — the vault's own retention statistic, unchanged.

V2-3 · The dipole keeps its address. OCEN-P398-PABH and OCEN-P397-P398 sit in the top decile of absolute net charge among eligible triangles.

V2-4 · The block map holds. The rank correlation between spent-era and sealed-era per-block drift across the fifteen populated charts exceeds the shuffled-block 95th percentile.

Horizon H2, wk1384-1539 — three sealed years; a typical triangle completes about three windings a year, so roughly nine windings make a legible behavior card.

V2-5 · The full pin map holds. The T4 behavior card recomputed on the sealed era correlates with the spent-record map across all eligible triangles, above the shuffled-address 95th percentile. And V2-1 is re-scored with the reversal band tightened to 14-22 — the three-year bootstrap sharpens to 15-20, and the promise grows teeth exactly as fast as the data can support them.

The post-hoc coda, registered now, no thresholds. After any Vault II scoring, the structures are also re-derived fresh on the sealed era — the same election recipe that produced the pool, the same grouping that produced the blocks, run on data that had no vote in the originals — and the memberships compared. Does never-seen ground re-elect the same witnesses? Redraw the same charts? This is labeled post hoc and exploratory: its job is not to pass or fail but to show whether the board’s organization reproduces. The full energy map runs post hoc on opened data the same way, labeled the same way. Frozen-by-name scoring gives the sharp verdicts; fresh-derivation comparison gives the portrait; the labels keep them from blurring.

The failure clause, stated up front. Any prediction that misses is reported as a miss in whatever future scoring we publish, and the registration itself is the deliverable either way. The last thing this paper hands you is a loaded copy of our experiment.

Appendices

A — The Player Cards. The key to reading a card, the deterministic recommendation rule box printed in full, the tier census, and six exemplar cards chosen by a printed rule — carried below. The complete 228-card deck (machine-generated per-station scorecards: individual metrics, group verdicts, percentiles-against-the-field, witness-phase state profiles, vault retention, and a deterministic recommendation with its rule printed) ships in the deposit crate as the canonical run file `PLAYER_CARDS_20260814_014430.md` (8 TRUST, 147 TRUST-WITH-NOTE, 39 UNRESOLVABLE, 33 WATCH, 1 CONTROL).

B — Vocabulary Traps and Prior Tools. The glossary: terms this corpus shares with other fields under different meanings, the false friends, and the prior-tool landscape with proper credit. Carried in full below.

C — Reproducibility and Provenance. Generated at deposit time by a MANIFEST run over the working tree, never typed by hand. The complete tool roster with checksums, the run-stamp registry, the vault protocol with its ratification and crash-rule record, and the one-double-click regeneration recipe. Everything in this paper regenerates from public data and pure-stdlib scripts; a result whose only home is prose is at risk, and none of these are.

D — The Station and Block Atlas. Built: `station_map v1.0`, run 20260818_144812, from the canonical clusters and station coordinates. All 229 stations coloured by their fifteen populated kinematic charts, with the inland-Oregon micro-block, the dipole stations and the witness pool labelled by name. Shown as Figure 1; the full-resolution vector and its per-block roster legend ship with the deposit.

The post-hoc analyses are summarized at §6/§7 and carried in full as §10, the Addendum and Epilogue, placed forward by design.

Appendix A — The Player Cards — per-station scorecards

Machine-generated by `PLAYER_CARDS v1.1` at 20260814_014430. ZERO recompute: every number is merged from an existing canonical run and names its stamp.

Sources found

- clusters: run_stage1_2026-08-11\clusters.json [stamp unstamped]
- scorecard: run_superstar_scorecard_20260813_212659\superstar_scorecard_data_20260813_212659.json [stamp 20260813_212659]
- station_phase: run_station_phase_20260812_183234\station_phase_data_20260812_183234.json [stamp 20260812_183234]
- vertex: run_vertex_screen_20260812_151555\vertex_screen_data_20260812_151555.json [stamp 20260812_151555]
- loop_sum: run_loop_sum_20260813_194104\loop_sum_data_20260813_194104.json [stamp 20260813_194104]
- steps: ABSENT — its fields render n/a
- witness: run_witness_phases_20260813_230157\witness_phases_data_20260813_230157.json [stamp 20260813_230157]
- vault: run_vault_open_20260814_000053\vault_open_data_20260814_000053.json [stamp 20260814_000053]

RECOMMENDATION RULES (deterministic, applied in this order; first match wins):

- R1 EXCLUDE-FOR-CALIBRATION - station on the stage-1 exclude/quarantine list
- R2 CONTROL - named positive-control station (PSPT)
- R3 UNRESOLVABLE - same-block census verdict UNRESOLVED (seam-dweller: needs a cross-chart method, not a verdict)
- R4 TRUST - frozen-pool member retaining n_pass >= 5 at the vault opening (stamp 20260814_000053)
- R5 TRUST-WITH-NOTE - frozen-pool member below the vault bar, or with unmeasurable criteria ('?')
- R6 WATCH - census CONVICT/SUSPECT or rogue-candidate flag
- R7 TRUST-WITH-NOTE - default: no adverse finding on record, not pool-vetted (the note says so)

Vocabulary: TRUST · TRUST-WITH-NOTE · WATCH · EXCLUDE-FOR-CALIBRATION · UNRESOLVABLE · CONTROL. Every number on a card names its run stamp.

Roster summary — 228 stations carded

- CONTROL: 1
- TRUST: 8
- TRUST-WITH-NOTE: 147
- UNRESOLVABLE: 39
- WATCH: 33

How to read a card

Every card is machine-merged from the canonical runs listed under Sources found, with zero recompute: each number carries the stamp of the run that produced it, and a field whose source run is absent renders n/a rather than being re-derived. The title line gives the station code, its block (the kinematic chart it belongs to in the deposited clusters), and its standing — FROZEN-POOL WITNESS (one of the 28 elected before the vault), CONTROL, or field station. **Individual** is the station alone: coverage (the fraction of record weeks it reports at all), gate (the fraction of its weeks audible past the amplitude gate), the scorecard's two activity metrics (med|W| and frac10) with its composite star score, the station's censused verdict (CLEAR / SUSPECT / CONVICT / UNRESOLVED) with the consistency behind it, and the mean-phase resultant R with its direction — R near 1 means the station winds with one preferred phase; n/a means too few phased weeks to say. **Group** is the same judgment at block level — what the station's whole chart says — because a station can only be fairly judged against neighbors sharing its ground truth. **Vs others** places the station in the field: percentiles for the activity metrics, the six-criteria screen (P pass, ! fail, ? unmeasurable) with its n_pass, the state profile from the witness-phase census (state codes with c consistency and r rate, and the count of state changes), and VAULT held-out n_pass — how many criteria the station retained at the one-shot vault opening. **RECOMMENDATION** is the rule box's first matching rule, printed with its rule number so the verdict is auditable, never editorial. The short field names are the tools' own; exact definitions live in the producing tools' headers, and every producing tool ships in the crate with its checksum in Appendix C's roster.

Six exemplar cards, chosen by rule

The full deck is 228 cards and lives in the deposit crate (below); the paper carries six, selected by a printed rule with no curation: for each of the five recommendation tiers present in the deck, the first card in deck order carrying that tier — GOLY (TRUST-WITH-NOTE, pool), P324 (TRUST), BAMF (UNRESOLVABLE), BCES (WATCH), PSPT (CONTROL) — plus the first card carrying the R7 field-default note, AL2H, because the default's wording ("no adverse finding on record, not pool-vetted") is itself a fence worth seeing on a real card.

GOLY · block 1 · FROZEN-POOL WITNESS

Individual: coverage 0.980 · gate 0.670 · med|W| 0.311 · frac10 0.282 [20260813_212659]; star 0.250 · census CLEAR (consist 0.670) · mean-phase R n/a / n/a deg [20260812_183234] **Group:** same-block census CLEAR · verdict n/a · shadow n/a [20260812_183234] **Vs others:** percentiles: med|W| p11 · frac10 p6 [20260813_212659]; six criteria: PPPPPP (n_pass 6) [20260813_212659]; state profile: PR c0.800/r0.591 DU c0.830/r1.310 DE c0.570/r0.575 · changes 2 [20260813_230157]; VAULT held-out n_pass 3 [20260814_000053] **RECOMMENDATION: TRUST-WITH-NOTE** — R5: frozen-pool witness, below the vault bar (n_pass 3).

P324 · block 12 · FROZEN-POOL WITNESS

Individual: coverage 1.000 · gate 0.650 · med|W| 0.323 · frac10 0.297 [20260813_212659]; star 0.480 · census CLEAR (consist 0.500) · mean-phase R 0.853 / -34.900 deg [20260812_183234] **Group:** same-block census CLEAR · verdict n/a · shadow n/a [20260812_183234] **Vs others:** percentiles: med|W| p12 · frac10 p20 [20260813_212659]; six criteria: PPPPPP (n_pass 6) [20260813_212659]; state profile: PR c0.500/r0.818 DU c0.500/r0.668 PO c0.670/r0.482 DE c0.500/r0.099 · changes 3 [20260813_230157]; VAULT held-out n_pass 5 [20260814_000053] **RECOMMENDATION: TRUST** — R4: frozen-pool witness, retained (n_pass 5) at the vault opening.

BAMF · block 1 · field station

Individual: coverage 1.000 · gate 0.940 · med|W| 1.054 · frac10 0.299 [20260813_212659]; star 1.520 · census SUSPECT (consist 0.800) · mean-phase R 0.624 / 39.200 deg [20260812_183234] **Group:** same-block census N/A · verdict UNRESOLVED (too few same-block partners) · shadow 1.520 [20260812_183234] **Vs others:** percentiles: med|W| p80 · frac10 p22 [20260813_212659]; six criteria: ?!PPPP (n_pass 4) [20260813_212659]; state profile: DU c0.600/r0.738 PO c0.600/r0.616 DE c0.800/r0.242 · changes 3 [20260813_230157] **RECOMMENDATION: UNRESOLVABLE** — R3: seam-dweller; fans cross chart boundaries.

BCES · block SB0 · field station

Individual: coverage 0.980 · gate 0.830 · med|W| 0.716 · frac10 0.357 [20260813_212659]; star 1.300 · census SUSPECT (consist 0.800) · mean-phase R 0.510 / 18.800 deg [20260812_183234] **Group:** same-block census SUSPECT · verdict MIXED — read the matrices · shadow n/a [20260812_183234] **Vs others:** percentiles: med|W| p55 · frac10 p93 [20260813_212659]; six criteria: !!!PPP (n_pass 3) [20260813_212659]; state profile: DU c0.600/r1.097 DE c0.600/r0.368 · changes 0 [20260813_230157] **RECOMMENDATION: WATCH** — R6: SUSPECT.

PSPT · block 11 · CONTROL

Individual: coverage 0.990 · gate 0.940 · med|W| 0.637 · frac10 0.299 [20260813_212659]; star 1.000 · census CLEAR (consist 0.670) · mean-phase R 0.175 / 45.100 deg [20260812_183234] **Group:** same-block census N/A · verdict n/a · shadow n/a [20260812_183234] **Vs others:** percentiles: med|W| p44 · frac10 p22 [20260813_212659]; six criteria: ?!PPPP (n_pass 4) [20260813_212659]; state profile: PR c0.500/r2.822 DU c0.800/r1.614 PO c0.500/r0.676 DE c0.670/r0.295 · changes 6 [20260813_230157] **RECOMMENDATION: CONTROL** — R2: named positive control; kept for instrument calibration.

AL2H · block SB0 · field station

Individual: coverage 0.260 · gate 0.480 · med|W| 0.351 · frac10 n/a [20260813_212659]; star 0.410 · census CLEAR (consist 0.670) · mean-phase R n/a / n/a deg [20260812_183234] **Group:** same-block census CLEAR · verdict n/a · shadow n/a [20260812_183234] **Vs others:** percentiles: med|W| p17 · frac10 pn/a [20260813_212659]; six criteria: PP?!? (n_pass 2) [20260813_212659]; state profile: DE c0.570/r1.257 · changes 0 [20260813_230157] **RECOMMENDATION: TRUST-WITH-NOTE** — R7: no adverse finding on record; not pool-vetted.

The full deck

Nothing was cut from this appendix; it was relocated. The complete 228-card deck is carried in the deposit crate as receipts_from_disk/PLAYER_CARDS_20260814_014430.md — the canonical run file itself, byte-for-byte, which this appendix's copy was verified identical to before the paper slimmed down to the key, the rule box, the census, and the six exemplars above.

Appendix B — The Same Object Under Other Names, and Who Built the Tools First

Bottom line up front: almost nothing in this paper is a new mathematical object. The triangle charge has been computed since 1988, the loop walk since 1958, and the holonomy family has a dozen names across as many fields. What this paper does that those fields do not is read the quantity they remove, and publish the reliability characterisation — the nulls, the gates, the registered failures — that instrument papers usually leave out. This appendix exists so that a specialist from any of those fields finds their own vocabulary waiting for them, credits their literature first, and sees exactly where our claim begins.

One object, many names

The costliest error class in this program was meeting an old friend under a new name and treating it as a stranger. The rows below are identities, not analogies — each row is one object.

this paper says	other fields say
triangle charge (a small closed loop returns an integer from wrapped phase differences)	residue (InSAR, optical holography, MRI phase unwrapping) · plaquette charge (lattice gauge theory, where the plaquette angle is defined mod 2π)
the loop walk (a directed phase sum around a closed triangle of stations)	closure phase (radio interferometry since 1958, through VLBI to the Event Horizon Telescope; ported to optical astronomy by aperture masking)
holonomy (what comes home changed after transport around a closed loop)	anholonomy · angular excess · Berry phase · Pancharatnam phase · Wilson loop · Hannay angle
winding	linking · circulation quantum · topological charge

Even our redundancy accounting is known art: for N stations there are roughly N^3 closure phases but at most $(N-1)(N-2)/2$ independent degrees of freedom, a bookkeeping the interferometry community settled long ago, and which this paper meets as discrete Stokes.

The credit ledger

Goldstein, Zebker and Werner introduced residues in 1988 (“Satellite radar interferometry: two-dimensional phase unwrapping,” *Radio Science* 23(4), 713–720, DOI 10.1029/RS023i004p00713): wrapped phase differences summed around a small closed path yield a charge of +1 or −1, branch cuts are drawn between them, and — their words — most residues come in plus-minus pairs. That is this paper’s triangle charge exactly, and the paired structure we call the dipole is their observation wearing our vocabulary; the phase-unwrapping world even ships dipole-remover preprocessors. The inversion is the honest boundary of our claim: they compute residues to remove them, because a residue is the obstruction that blocks unwrapping an image. We compute them to read them. Their signal is our noise; our signal is their nuisance.

Jennison introduced the closure phase in 1958 (*MNRAS*, DOI 10.1093/mnras/118.3.276; extended by Rogers and colleagues in 1974) for a reason this paper should salute: phase errors on an antenna array are antenna-based, and they cancel in the directed sum around a triangle. That is the unreliable-sensor problem solved topologically, sixty-eight years ago. The contrast with our use is the same inversion in a different key: their loop cancels the station error so the sky can be imaged; our census walks the loops to name the station. Closure cancels; the census names.

Interferometry has since arrived at the gauge language on its own: Thyagarajan, Nityananda and Samuel frame closure quantities as invariants of an Abelian gauge theory (*Phys. Rev. D* 105, 043019, 2022), with an active modern lineage in closure statistics (Blackburn et al. 2020, *ApJ* 894:31; Broderick and Pesce 2020, *ApJ* 904:126). The field already speaks holonomy without the loop-walker framing. All three anchor citations above were verified against the publisher record on 2026-08-14; anything added to this appendix later gets the same check, because a citation quoted from memory is a claim without a measurement.

Traps this paper defuses in advance

The triangle charge is winding mod 3. A single triangle’s wrapped-phase circulation can only report −1, 0 or +1: a true winding of +2 across one triangle reads as −1, and a true +3 reads as 0 — we measured both on synthetic loops. So no sentence in this paper equates the triangle charge with the winding number unhedged. The redemption arrives in the same breath: summing triangles into an n -gon reads winding mod n , so large loops resolve what small ones alias — that is discrete Stokes doing anti-aliasing — and the physical fence is that a true winding of magnitude 2 inside one ~40 km triangle would require the residual vector to turn twice along that small perimeter at millimetre amplitudes, which we document as a ceiling rather than fear as a live wound. Lattice gauge theory has always defined the plaquette angle mod 2π ; this is known art, cited as such.

Three words in this paper carry one defining sentence each, stated here so no specialist reads their own field's meaning into ours. Residual means the detrended, common-mode-removed displacement — not observed-minus-predicted in the regression sense. Ternary means the sign of a reading, taking values in -1 , 0 , $+1$ — not base-3 arithmetic and not three-valued logic. And a zero is only a measurement when the triangle was defined that week: an undefined triangle-week is an absence, and this deposit's own grouping defect (\$1, \$7.2) is a monument to what happens when an absence is allowed to impersonate a value.

What we do not claim

This instrument is not a better interferometer, a better unwrapper, or a better anything in the fields above, and no sentence here should be read as a challenge to their tooling. What we publish that instrument papers usually do not: the null bands measured before the questions were asked, the integrity gates that refuse to write when the record does not reproduce, the registered designs reported including their failures, and a one-shot vault opened exactly once. The contribution is the reliability characterisation of a deliberately simple channel — and the standing habit, paid for repeatedly in this very deposit, of checking the artifact rather than the intention.

Appendix C — Reproducibility and Provenance

The paper promises this roster twice, so here it is rather than a pointer to it. Every tool that produced a number in this paper, every launcher that runs it, the frozen registration documents the text cites by checksum, and the two interactive supplements — each with the first twelve hex digits of its SHA-256, read from the deposited file itself. A reader who wants to know whether the file in their hands is the file this paper ran can hash it and compare.

file	bytes	sha-256 (first 12)
tools/BLIND_RING.py	33,897	8ac947ac4fb6
tools/CASCADIA_CENSUS.py	32,700	aecf680def9c
tools/COHERENCE_TEST.py	28,152	d9cada2aa217
tools/MANIFEST.py	12,173	b919fe4afb7b
tools/PACHINKO_SUITE.py	58,583	c0a0f9074f91
tools/PRECURSOR_POSTHOC_V2.py	65,383	cf7c93dc98b2
tools/RUN_BLIND_RING.bat	45	c9d6d6198f47
tools/RUN_BUILD_PAPER.bat	864	e595b3206b77
tools/RUN_CASCADIA_CENSUS.bat	50	1db37a5cf5ea
tools/RUN_COHERENCE_TEST.bat	49	a55e26265997
tools/RUN_LINT_DEPOSIT.bat	1,348	0523a101dfde
tools/RUN_MANIFEST.bat	43	6d99dc6e325b
tools/RUN_PACHINKO_SUITE.bat	48	cc87ab3f360c
tools/RUN_PRECURSOR_POSTHOC_V2.bat	55	5355c52fe45f
tools/RUN_SEAM_WALK.bat	44	47a9c3741c71
tools/RUN_STATION_RING.bat	47	e4ae350ecf5e
tools/RUN_TREMOR_DRIFT.bat	47	ccc8fc034e34
tools/RUN_WATERFALL_MAP.bat	47	f65d2630a1d6
tools/RUN_WINDING_BUDGET.bat	48	a20fdffdacbc6
tools/RUN_WINDING_RHYTHM.bat	49	4c1f86d7cf82
tools/SEAM_WALK.py	39,370	abfa60669dfe
tools/STATION_RING.py	36,238	ef0fefd40869
tools/TREMOR_DRIFT.py	34,643	3a99e0ea7260
tools/WATERFALL_MAP.py	37,129	88dbc806a7b8
tools/WINDING_BUDGET.py	26,669	efcc4a86e330
tools/WINDING_RHYTHM.py	35,738	ea7b01f480b8
tools/build_paper.py	11,188	f24ef45c28d3
tools/ets_dates.json	72,148	c7901eaa3d11
tools/lint_deposit.py	8,467	488756b68878
registration/PACHINKO_PROTOCOL_v0-3_2026-08-21.md	15,416	1a5286f4de6f
registration/PACHINKO_PROTOCOL_v0-4_2026-08-21.md	18,318	fb8c33ca8188
supplements/BLOCK_EXPLORER_20260819_224258.html	17,464	e668a49ceecd
supplements/LIGHTNING_MAP_20260819_224304.html	488,373	9316cc2e00fe
supplements/RUN_BLOCK_EXPLORER.bat	52	b7cee68a112f
supplements/RUN_LIGHTNING_MAP.bat	51	24ac9c86d33f
supplements/block_explorer.py	17,529	a4d8c1886eea
supplements/lightning_map.py	33,634	008b18cedc2e

The run-stamp registry. The runs of record live in `receipts_from_disk/`, each in its own `run_<tool>_<stamp>/` folder, and every output file carries its tool's name and run stamp in the filename so it keeps its identity if it is ever moved out of that folder. The paper cites each of them by stamp at the point of use, which is the registry: a number

and the run that produced it never appear apart.

The regeneration recipe. Each tool is one double-click on its .bat and writes its own stamped run folder. RUN_BUILD_PAPER.bat rebuilds this PDF from the markdown beside it; RUN_LINT_DEPOSIT.bat runs the pre-deposit check over the result. Everything here regenerates from public data and pure-stdlib scripts. A result whose only home is prose is at risk, and none of these are.