

The Tower Is Coordinate-Free

What Survives When You Change the Vertical Axis of a Prime Walk

Matthew “Bear” Schulz · Independent Researcher · Belmont, CA ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com *Preprint, not peer reviewed. CC BY 4.0. Draft v3.4, 2026-08-22 (v3.3 deposited 2026-08-01). Companion to Taking a Loop Walk Through the Integer Forest (10.5281/zenodo.21449569).*

BOTTOM LINE UP FRONT

What this is. Take any whole number. Break it into its prime factors. Now walk that factorization like a set of compass directions — one leg per prime, a turn of 137.5° between legs, a climb each time you meet a new prime. Every integer gets a unique address in space. That is the construction from the companion paper.

The question this paper asks. The picture has a vertical axis, and the choice of what to put on it was mine. So: **which of the things we measured actually depend on that choice, and which would hold no matter what I had picked?**

The answer. Eight of the ten measured regularities never touch the vertical at all. They are computed entirely from where the walk goes sideways, and climbs never move you sideways. **They are identical under any vertical you choose.** One regularity reads the vertical, and it survives only under the original choice.

That is the whole paper, and the result is a negative one in the best sense: most of what we found is not an artifact of how I drew it.

Three things fall out along the way.

1. **A closure deficit with a floor you cannot argue your way past.** Every integer’s walk fails to close by some amount. We show that amount can never be driven to zero except in the exact cases where it is zero — and the reason is a theorem from 1966 about logarithms. This is the only substrate in this whole program where the noise floor is *proved* instead of measured.
2. **The leftover is Benford’s law.** We predicted the residual would be evenly spread. It isn’t. It follows the same distribution that governs leading digits in real-world data, and that turns out to be exactly the right answer for reasons the paper explains.
3. **The oldest example of this is musical.** Twelve perfect fifths do not close into seven octaves. They miss by 23.46 cents, everyone who has ever tuned a piano has dealt with it, and it is the same object.

What we are not claiming. No new theorems. Every core object here is classical, and where we found our “new” thing already published we say so and cite it. See *Standing Beside the Prior Work* near the end.

SECTION 0 — HOW A GUY WHO IS BAD AT MATH ENDED UP HERE

(Author’s voice. Provenance, not proof. Skip it if you want the results.)

When you are handed a hammer, everything starts looking like a nail. When you are handed a tool that can wear all sorts of cross-domain splice kits, and you walk enough loops, you start recognizing patterns.

I live in a block universe. Every loop closes.

I left math behind a long time ago because I struggled with a language that had concepts in it that felt wrong to me. Math was my fickle friend at best and my fiercest adversary on most days, and my grades show it. A lot of people can relate to hitting a wall with something new — that feeling that it just won’t stick. That was math for me. Weirdly, only certain parts blinded me. I don’t have dyslexia. I could read and perceive numbers fine. But the way it was taught didn’t click. I could spec-dump before a test, manage through the memorized concepts, pass — often barely — and then watch it disappear weeks later. That was calculus. When calculus fell, it took down a lot of career paths with it. At an engineering school, that is most of mechanical and electrical.

I could say I am just math-stupid, but that isn’t quite right either. I goofed off, played computer games, played sports, hated to study. And my math deficit had a weird smart spot. Statistics. Stochastic and

deterministic processes, Poisson, linear optimization, regressions, matrix algebra, discrete dynamical systems — those welcomed me. I think it clicked because it matched something already true in my head: statistics are frighteningly accurate across a population and break down completely at the individual.

So when I say I took a different path here, I mean I am an engineer who struggles with math.

A note on words, because words carry agency. This work gets called “linear and log towers,” and that phrasing messes people up — it messed me up. For me it is **Log and Anti-Log**. I am putting the word “exponent” in the swear jar. Same with “imaginary” numbers, which were horribly mislabeled or maybe just badly taught. A lot of people memorize the name and never learn the concept. For me it stayed a swear word for decades, through nearly every teaching method I could find — all of it ugly and depressing in a way the concept underneath never was. What finally cracked it was history and etymology, and by then the joke was on me. I had been mapping with quaternions from the start, so I was already living with the thing while still swearing at the name; once I saw that, making peace was not optional. Quaternions compose rotations, which is exactly what they are for, and this whole construction is a sequence of rotations.

And I have a grievance with zero.

If zero is a fairly recent invention, historically, what did people do before it — and why are some zeros true and others abstract? The universe I see is rates, and changes of rates, and changes of changes of rates. That is why I say the universe is ternary, and it is not because a sine wave goes +1, crosses zero, −1. It is because **the slope is the rate of change**, and the oscillating curve is the rate of *that*, and the interesting thing is the crossover — the instantaneous moment where the tangent grazes and the slope reverses. A regulator. Always predictable. That is the ternary I care about, and in it the idea of zero mostly falls away.

Where this actually started: a poker table. We were building adversarial signal processing for a dealer watching a loop of wagers close around a hand. Each player is a Byzantine sensor. Information gets repriced and bluffed. Follow the cards around the table and you have traced a loop. Connect the first and last point and you have enclosed an area. That area, and the record of how you got around it, is the voyage of the object through space and event-driven time — cards dealt, turns taken, initiative exercised to make the other sensors believe a false price.

From there it was the same shape everywhere I looked. Standing waves in a crystal family. Arrays that ring like a bell. Stars, quasars, the junction seams of tectonic plates. The tool was good at a lot of things.

Then I took a break to catch up on my math. Concepts only — and any time I say I am learning math or catching up on math, that is what I mean. Concepts I need to understand. Me asking all the stupid questions without touching a calculator. So I decided I was going to actually teach myself Chern-Simons instead of goofing at it. And once you are weaving looms around a torus and watching how eigenstates behave at modulus zero and modulus one — unit circles, unit spheres, unit areas, unit everything — you start quantizing the universe by choosing unit values for different geometries. You see how squares and n-gons and circles each hide their curvature somewhere different, and you start building a rulebook for uncovering it.

Deep in that, chasing repulsion and electron shells and why a nearly-full shell is the one that can pass traffic, the trail led to the Riemann Hypothesis.

RH doesn’t have the dignity of being a cemetery. It’s a graveyard. A Boot Hill for everyone who trades a day of math to learn why unsolved things stay unsolved.

I never got anywhere near solving it and I am not trying to. But the negative space around it screams for shape, and RH handed me things I could use: rungs, trivial zeros, base 2 and base 3, repulsion and steady-state eigenstates. It is also where my grievance with zero got sharpened. **The zeta zeros are true zeros — they earn the line.** The trivial zeros are fakers. They never earn a spot; they are the universe saying *eventually this is as good as zero*. My block universe says that isn’t closure.

This paper is not about RH and I want to be clear about that. But the log spiral in RH kept bothering me until I built something that answered it inside my own head, and this tower is what that turned into.

What the tower gives me is an energy network with two ways to reach the same address. A long walk, one step at a time. And a short walk — a straight hypotenuse with a magnitude, from origin to a unique address. The gap between them is the closure deficit, and the whole rest of this paper is about that gap.

It also gave me the subway. Riding the anti-logs of primes is a free ride along the cylinder wall: the wall has circumference = 1, so a full turn costs nothing at all. A free subway ride on the circumference-1 tower network. And trivial zeros became more rungs.

This has been about three months of work, mostly in bursts. I don't want to oversell that. If my math-stupid brain got helped by packing the concepts into a picture I can walk through, maybe someone else finds it useful too. Is it finished? I hope not. It has a lot left to teach.

For those of you who stare at numbers and wonder — here is my attempt at artwork. Well, maybe more a fingerpainting hiding in plain sight. It's here to amuse you, and if in that moment something clicks and you say to yourself *I thought so* — that's the win. That's the whole thing. A little map and a shortcut, from 1, 2, 3 out to a hundred billion. The hypotenuse. The 3D dead-reckoning collapse back to the origin.




Figure 1. The forest resolves with count. Side view, rung on a log scale, at $N = 1,000 \cdot 10,000 \cdot 100,000 \cdot 1,000,000$. The same construction at four sample sizes: the structure sharpens rather than changes, which is the first hint that what the picture shows is in the integers and not in the drawing.

Origin note, added in v3.4 — where the bridge's name came from. The exp/log bridge of this paper has a provenance worth recording honestly. In April 2026, early in this program, an autocorrect changed an intended reference to *Sophie's World* into “Sophomore’s Dream,” and the substitution went unnoticed long enough to be taken seriously. Followed deliberately, the accident pointed somewhere real: the Sophomore’s Dream identities (Johann Bernoulli, 1697) — $\int_0^1 x^{-x} dx = \sum n^{-n}$ and its alternating twin — are the exp/log gearbox in miniature, a smooth integral becoming an integer-indexed sum through the expansion of $\exp(x \cdot \log x)$. That is exactly the shape of the two-tower map measured in the energy-bridge section below, where $E = \log n$ is the additive tower, $n = \exp(E)$ the multiplicative one, and $-F'(s)$ is the gear that meshes them. Nothing in that section depends on the anecdote — the identities are classical and the bridge is classical calculus in the tower’s coordinate — but the pointer arrived by accident, and part of this program’s method is saying plainly where its instruments came from. The bridge was found because an autocorrect was taken seriously enough to test.

Changed in v3.4 (2026-08-22): the arithmetic-metric citation now names its author (D. Dominici) in both places; this origin note was added; the interactive bench in the supplement was updated to v1.3 (a crash fix and two corrected readout labels; no construction change); and the Zenodo identifiers for the companion and related deposits were moved to concept DOIs, which always resolve to the latest, corrected version of each work. A prior-art sweep (2026-08-22) also strengthened the floor’s citations — the operative effective exponent for $a \cdot \log 2 - b \cdot \log 3$ is Rhin’s ≈ 7.616 (1987; Wu 2003; Zudilin 2004), far sharper than the general Matveev route previously cited alone — and added the standard references for the Benford dictionary, the chain count, and the divides-relation graph. A catalog sweep did not locate the sideways constant. Two late additions close the record honestly: a one-paragraph result in the subway section — the Liouville sign as the holonomy of a flat $\mathbb{Z}/2$ connection, $\lambda(a)\lambda(b) = (-1)^{\{d(a,b)\}}$, verified and offered at observation weight — and a Threads Left Deliberately in the Open section that hands off the uncomputed diffraction pattern of the prime-rank point set (to the aperiodic-photonics program of the Dal Negro group), the analysis of the sideways constant, and one fenced Hofstadter rhyme. And the boundary-holonomy question the first edition stated as open was run rather than handed off: the transition vertices carry exactly zero vertical vector area — no boundary holonomy in the measured channel — and 99.92 percent of the horizontal area. One prior open item inside the paper is now a measurement. No previously published value changed.

THE CONSTRUCTION, AND THE TWO THINGS BOTH CALLED “THE HYPOTENUSE”

The walk moves **sideways** when it walks a prime’s exponents — heading $\text{rung}(p) \times 137.50776^\circ$, one unit step per exponent. It moves **up** when it meets a new prime. Sideways moves never change your height; climbs never change your sideways position. That clean separation is what the whole paper turns on.

Here is a hygiene problem the companion paper left implicit, and it cost me a day. The construction carries two different quantities and both got called “the hypotenuse”:

name	what it is	what uses it
the flat chord $C(n) = \sqrt{x^2+y^2}$	straight-line distance home, ignoring height	the closure deficit
the full reach $R(n) = \sqrt{x^2+y^2+rung^2}$	straight-line distance home in 3D	the octave law

Measured: $C(127) = 1.75547$ and $R(113) = 30.00110$. Same construction, different questions, different answers. **Name which one you mean in the same sentence as the claim, every time.** The rest of this paper does.

A correction owed to the published companion. Its construction text says the walk climbs “to vertical height equal to the running count of distinct primes used.” That is not what the code does, and it is not what reproduces the numbers that paper printed. The climb is the **rung of the largest prime factor** — the index of the biggest prime in the factorization, not a tally of how many primes there are. Only that reading returns $R(113) = 30.00110$ and $R(100) = 4.26423$; the running-count reading returns 1.03 and 3.63 instead. So the companion’s numbers were right the whole time and one sentence of its prose was not. An erratum will carry the fix.




Figure 2. The Integer Forest at $N = 10,000$. Orthographic composite view, azimuth 45° , elevation 45° , rung on a log scale with $3\times$ vertical exaggeration. Every point is one integer at its own address; the flared skirt is the many integers of low rung, and the column is the run of large prime factors.

THREE WAYS TO READ ONE INTEGER

Every integer is a list of prime exponents — a point in a lattice. The tower reads that list three ways:

- **Which prime is biggest** (its rung) — this is *roughness*.
- **How many prime factors total**, $\Omega(n)$ — this is *depth*, and it is exactly the walk’s length minus the stem.
- **The energy**, $E(n) = \log n$ — the list dotted with the log of each prime.

The energy is **perfectly additive**, with zero left over, because log turns multiplication into addition. $E(165) = \log 3 + \log 5 + \log 11 = 5.10595 = \log 165$, exact. Each prime carries a fixed quantum of energy; every composite is the exact sum of its parts.

A climb is a ratio, not a big number. $\Delta E = \log(b/a)$. The log is only there to make ratios add instead of multiply, which is why the energy adds at all. That is not decoration — the next two sections turn on it.

And there is a fourth reading, new here: the turning number. Close the walk with its chord and count how many net times the direction of travel rotates. It is an integer for every n , to 4×10^{-16} across $n = 2 \dots 20,000$. It is *not* a function of how many primes you have — at three distinct primes it spreads across $\{-2, -1, 0, 1, 2\}$ — because it reads the *gaps between rungs*. And it is completely blind to exponents: 12, 24, 48, 96, 192 all return -1 while Ω climbs from 3 to 7, because repeated primes are collinear legs and a straight-through corner turns nothing. It reads the *set* of primes and nothing about multiplicity.

(Caveat, because it matters elsewhere: the turning number is a planar quantity. Project a 3D curve two different ways and you get two different answers. It is bulletproof on the tower, where the projection is the construction, and it is a poor invariant for anything embedded in three dimensions.)

WHAT SURVIVES A CHANGE OF VERTICAL — THE MAIN RESULT

regularity	what it reads	verdict under any new vertical
distinct integers get distinct addresses	position	invariant
φ^2 rungs per revolution = 2.618034	angle	invariant
walk length = $1 + \Omega(n)$	leg count	invariant
ancestry: grazing an axis means divisibility (13,197 events, 0 unexplained)	sideways only	invariant, identical
stiff spreading, exponent 0.977	sideways	invariant
frame-lock ceiling 1.072924, tune points at Fibonacci	rotation	invariant
Fibonacci returns at Δ rung 1,2,3,5,8...	angle	invariant
Liouville parity excess 531	pure arithmetic	invariant
the octave law, $k = \text{rung} + 1$	full reach	only under rung
the $p = 5$ anomaly	full reach	rides the octave law

The first eight are the paper’s central prediction and they reproduce **identically** — not similarly, identically — because they never touch the vertical. That is empirical proof the results are not an artifact of the drawing.

Render the same data under four different verticals and only one produces the tower. Height = rung gives the prime-factorization structure. Height = Ω collapses to about fifteen flat bands. Height = $\log n$ smears into a wedge. Height = n gives a featureless column.

THE WALLS ARE FLAT, AND WE MEASURED IT

This is the “curl” question, and the answer is zero.

Treat the closed loop — walk out, chord home — as a circuit and compute its vector area, $S = \frac{1}{2} \oint \mathbf{r} \times d\mathbf{r}$. That is the Stokes object whose flux through any named field would be a holonomy. It is the curl-like quantity, and it is the right thing to ask for.

Measured: along the cylinder walls the transport carries no holonomy of its own. The golden-angle offsets all share the same vertical axis, so they *commute*. Commuting transport is abelian, and abelian transport is flat. The walls are curvature-free by measurement, not by assumption.

That is why the construction’s only closure defect lives in the **loop as a whole** and not in any wall you

walk along. The vertical component of the loop's vector area is the decomposed height-residual of the walk, and for a prime the area is a tall thin sliver of base 1 and height equal to the climb, so its magnitude tracks the rung.




Figure 3. Anatomy of the construction. Each cylinder's axis sits *on* its cofactor's cylinder, so the nesting chain is the factorization itself; the radius is the top exponent, and a walk grazes exactly its ancestor axes — a maximal chain in the divisor lattice. The walls measured flat above are these walls; the caps and bases are the boundaries named as untested in the next paragraph.

What is still open, and stated as open. The walls are flat. The *boundaries* are not tested. Where a walk enters or leaves a cylinder — the caps and bases — is exactly where a new heading first appears, and it is the one place in the construction where the commuting argument does not obviously apply. **Whether transport across a cylinder boundary carries holonomy was unmeasured in the first edition; v3.4 runs the measurement two paragraphs below.**

It is also well posed and probably cheap: compare the vector-area contribution at the transition vertices against the contribution along the legs. We have not run it, we are not claiming an answer, and we would rather someone else got there first than have this sit as a private hunch. **That is the next measurement in this construction and it is stated here so it can be taken.**

Taken — in v3.4, by this construction, the day the question was re-read. The circuit integral decomposes exactly by segment: on a straight segment from a to b the contribution to $S = \frac{1}{2} \oint \mathbf{r} \times d\mathbf{r}$ is $\frac{1}{2}(a \times b)$, so every unit of vector area can be charged to a named piece of the loop — the stem, the climbs at the transition vertices, the horizontal legs, and the hypotenuse home. Run over all integers to $N = 100,000$, the ledger closes cleanly. The stem and the hypotenuse contribute exactly nothing: both are radial segments, aimed at or from the origin, and a radial segment carries no moment. The climbs — the boundary crossings the question asks about — contribute exactly zero to the vertical component: a climb's displacement is purely vertical, so its cross-product has no vertical part, structurally, and the run returns 0.000000. The vertical vector area — the closure-deficit channel this paper measures everywhere else — is therefore carried by the legs alone, which means the flatness of the walls extends to the boundaries: **transport across a cylinder boundary carries no holonomy in the measured channel.** What the boundaries own instead is the horizontal area: the climbs carry 99.92 percent of $|S_{xy}|$ against 0.08 percent for the legs — those are the tall thin slivers named above, and they turn out to be nearly the whole horizontal story. A clean separation by channel: vertical to the legs, horizontal to the climbs, nothing to the radial ends. Stated at measurement weight — one scale, one session, exact per-segment

decomposition — and it regenerates in thirty lines from the construction in the supplement.

THE ENERGY BRIDGE, AND WHY IT IS A LAPLACE TRANSFORM

Write $\lambda(n) = (-1)^{\Omega(n)}$, the Liouville function — plus one if n has an even number of prime factors, minus one if odd. Sum it with a weight and you get $F(s) = \sum \lambda(n)/n^s = \zeta(2s)/\zeta(s)$, a classical object.

Differentiate and you get $\sum \lambda(n) \cdot \log(n)/n^s = -F'(s)$, confirmed numerically: $-F'(2) = -0.291233$ and $-F'(3) = -0.118112$, matched to six digits.

The energy axis rides that exactly. The sideways golden-angle chord does not.

Why the energy axis and not the other one

The companion paper stated this and left it as a bare fact. It has a reason, and the reason deserves more than the one paragraph I originally fenced it to.

A Dirichlet series is a Laplace transform. Write it out:

$$\sum \lambda(n)/n^s = \sum \lambda(n) \cdot e^{(-s \cdot \log n)}$$

That is the Laplace transform of the Liouville sequence, sampled at times $t_n = \log n$. The tower's energy axis is the transform's time variable. That is not an analogy and not a renaming — it is what the object is.

Once you see that, the result stops being a coincidence. Differentiating a Laplace transform with respect to s pulls down a factor of $-t$, and here t is $\log n$, which is the energy. **So $-F'(s)$ is the energy-weighted sum by construction.** The bridge lands on the energy axis because the energy axis is where the transform's clock lives. Nothing else in the tower is that clock, which is why nothing else rides it.

That also explains the higher derivatives cleanly. Differentiate k times and you pull down $(-\log n)^k$, so energy ^{k} sums equal $(-1)^k F^{(k)}(s)$ — verified to twelve digits at $s = 3$.

Where the geometry ties in, and where it doesn't. The tower's *reach* connects to the derivatives cleanly at exactly one place, $k = 2$, and it does so by plain Pythagoras: $\sum \lambda R^2/n^s = F''(s) + C_2(s)$, because $R^2 = C^2 + \text{rung}^2$ and the sum splits. At $k = 1$ and $k = 3$ the reach is energy ^{k} plus a sideways correction with no closed form. **The geometry crosses the bridge once, at the one power where the metric is additive.** That is a satisfying answer to why it works at all.

The sideways constant. The chord sum converges to -0.32774696 at $s = 2$ and -0.13674984 at $s = 3$. Running `mp.identify` against the standard basis returns no closed form. We assert it only as *not located*, and the caveat in *Standing Beside the Prior Work* applies with force. A literature and catalog sweep (2026-08-22: OEIS and web-indexed constants; exponential-sum literature with quadratic-irrational rotation) did not locate this constant; “not located” remains the ceiling of the claim.

What this paper does not do. $F(s) = \zeta(2s)/\zeta(s)$ has poles wherever $\zeta(s)$ has zeros, so the Liouville series is formally tied to the zeros of the zeta function. That tie is real and it is classical, and nothing in this paper bears on it. The tower uses the Dirichlet series as bookkeeping for the energy axis and stops there. Nothing measured here is evidence about the zeros, and reading it that way would be a mistake.

THE SUBWAY — WHY SHARED FACTORS ARE A FREE RIDE

If going around a cylinder costs nothing, then the only thing you pay for is changing exponents. That gives a distance between two integers:

$$\text{subway}(a, b) = \Omega(\text{lcm}(a, b)) / \gcd(a, b)$$

Going the long way through the origin costs $\Omega(a) + \Omega(b)$. So the interchange saves exactly $2 \times \Omega(\gcd(a, b))$ — the shared factorization you never have to unbuild. $25 \rightarrow 165$ share a single 5, so you pay 3 instead of 5.

This is an algebraic identity, not a lucky observation. **The distance formula is prior art:** D. Dominici, *An Arithmetic Metric*, arXiv 0906.0632, gives the identical expression as the ℓ_1 taxicab distance on the exponent lattice. We cite it and we do not claim the formula.

But the formula is a shadow of the tower, not the other way round, and the comparison should be made honestly in both directions.

The taxicab metric on the exponent lattice returns exactly one thing: a number $d(a, b)$. It has no positions, no geometry, and nothing to look at. Every integer is an abstract lattice point and the metric tells you how far apart two of them are.

The tower returns that same number **and**:

	in the taxicab metric	in the tower
distance between two integers	yes	yes — identical
a position for each integer	no	yes, unique, injective to 10^6
a chord from origin to that position	no	yes, $C(n)$
a closure deficit between the two routes	no	yes, and it is the paper’s central quantity
cylinders whose radius is the top exponent	no	yes, exactly, to 10^{-15}
grazing = divisibility (the ancestry theorem)	no	yes, 13,197 events, zero unexplained
the octave bracket $2^{(\text{rung}+1)}$	no	yes
a turning number per integer	no	yes

So the honest statement runs the other way round. The formula and our construction are different objects with different reach. Someone wrote down the distance; we built a space that has that distance in it, along with seven other things the distance cannot express.

And the space was not built to compute distances. It was built to run a loop walk on a substrate where the answers are already known — pure quantized winding structures living inside the arithmetic itself — so the instrument could be calibrated against them before it was pointed at anything softer. Most instruments of this kind are domain-specific: one tool, one substrate. This one is closer to the whole knife. Seeing the forest is the point, and the distance formula is one tree standing in it. The citation is owed either way.

The same care applies to the factorization-on-a-torus prior art (Lulek & Olchawa 1996). They realised factorization as motion on a torus with one perpendicular axis per distinct prime. That is genuinely the same core idea and we do not claim it. It is also not this construction, whose headings are golden-angle offsets indexed by prime **rank** rather than perpendicular axes indexed by prime identity — which is what produces the cylinders, the grazing relation, and the non-closing helix. **Cite the ancestor, scope the claim, and do not hand over the whole map because someone drew one of the roads.**

How far the free ride goes. 2^{80} and 2^{81} are 1.2×10^{24} apart on the number line and **one hop** apart on the interchange. From 2 to a billion-scale prime is two hops. The gap on the number line diverges without limit while the interchange distance stays fixed.

The long walk versus the short walk, as a ratio. Reaching n by unit steps costs $n - 1$; reaching it by interchange hops costs $\Omega(n)$. The ratio grows like $n / \log n$ — mean 2,033 over the first ten thousand integers, and its maximum always lands on a prime, because one hop reaches what the long walk pays $n-1$ for. **Roughness is cheap; walking is exponentially expensive.**

A small result the subway was carrying all along, added in v3.4. The subway distance and the Liouville function are the same object seen twice. Because Ω is completely additive and $\text{lcm} \cdot \text{gcd} = ab$, the metric expands as $d(a,b) = \Omega(a) + \Omega(b) - 2 \cdot \Omega(\text{gcd}(a,b))$ — and the last term is even, so the parity of the distance never depends on what a and b share. Read that as signs and it says $\lambda(a) \cdot \lambda(b) = (-1)^{\{d(a,b)\}}$: ride the subway from a to b , count the stations, and the parity of your fare is the product of the two Liouville signs, no matter which route you take. We verified it the boring way — four thousand random pairs up to 20,000, zero violations, and path-independence checked explicitly across five different routes from 25 to 165. In the language of geometry this is a flat $\mathbb{Z}/2$ connection on the integer network whose holonomy is the Liouville function: transport a sign around any loop and it comes back unchanged, because the network’s curvature is zero — the whole of λ is stored in the endpoints, none of it in the path. Both halves are classical (complete additivity of Ω ; $\lambda = (-1)^\Omega$ is Hardy & Wright material); a 2026-08-22 literature sweep did not locate the two welded together as a connection-and-holonomy statement, so the welding is offered here at exactly that weight: a one-paragraph observation on classical ground, free to anyone who can build on it.

The arteries

Which primes are the trunk lines is a clean law on two axes. **Usage** — the fraction of integers that use prime p — is exactly $1/p$: 2 at 50%, 3 at 33%, 5 at 20%, 7 at 14%. **Transfer** — the chance p is the shared factor in a move — is $1/p^2$: 25%, 11%, 4%, 2%.

Ranking is strictly $2 > 3 > 5 > 7$ on both. So **2 owns usage** and the octave lane, **3 is the quiet number**

two, and **5 owns geometry** — it is the one prime that breaks the octave law, and that is a different axis entirely.

THE OCTAVE LAW, AND THE ONE CRACK THAT SHOWS UP THREE TIMES

A prime at rung r sits at full reach $\approx r$, and the smallest power of two that exceeds it is $2^{(r+1)}$. The bracket closes within **100/r percent** and tightens as you climb: 33.3% at rung 3, 3.3% at rung 30, 0.60% at rung 168, 0.33% at rung 304.

The asymmetry is worth sitting with. To match a rung-30 prime by a power of two you need $2^{31} \approx 2.1$ billion — nineteen million times the magnitude, to buy one short step at that height.

And there is an $n-1$ law, which the companion paper filed as “a pull, not a proof.” It is now measured, in two forms.

Every prime’s reach sits strictly between its $(r-1)$ and $(r+1)$ power-of-two reaches. Checked at 13, 113, 997, 2003 — all bracketed cleanly. **The one exception is $p = 5$** , where $\text{reach}(5) = 3.495$ slips just past $\text{reach}(2^4) = 3.479$.

So the prime resolves in the $(n-1, n+1)$ window for every prime **except the same 5 that breaks the octave law**. The $n \pm 1$ law and the $p = 5$ anomaly are one exception, not two. A construction with two unrelated flaws would not put them on the same integer.

Second form. The tower’s rotation number is 0.381966..., whose continued fraction is $[0, 2, 1, 1, 1, 1, \dots]$ — all ones, which is the definition of the golden mean, the most irrational number there is. Its Fibonacci convergents bracket it with perfectly alternating sign, $n-1$ below and $n+1$ above, closing in from both sides.

A null, reported honestly. The same overtake measured against the *energy* reach has no stable percentage — it drifts between -19% and -36% , because the power of two’s sideways walk swamps its energy reach. The clean closure law is a rung-tower fact and nothing more.

THE FLOOR IS A THEOREM — THE PAPER’S STRONGEST RESULT

Why zero has to be isolated for any of this to mean anything

Every instrument in this program reads the same shape: an integer channel plus a continuous leftover, where **leftover zero means the loop closed**.

That reading only works if *closed* can be told apart from *very nearly closed*. If the leftover could be squeezed arbitrarily close to zero without ever reaching it, then closure would be a limit and not a measurement, and the integer channel would be reading fog.

Geometrically, zero curvature is a knife edge — a boundary between two signed regimes, with nothing on it. But a *reading* of zero is a band straddling that edge, as wide as your noise. **The truth is a razor; the measurement is a smudge across it.** Every substrate this program has touched answers that by measuring the smudge: 35.81° for protein defect angle, 0.1886 for protein writhe, 0.0640 for RNA, $1.360^\circ/\text{bp}$ for DNA twist. Each cost a run. Each was substrate-specific. One of them, once finally measured, killed a claim we had already written down.

On the integers, you don’t measure it. You cite it.

Ask the sharpest version: can two prime lines close on each other? How small can $|a \cdot \log 2 - b \cdot \log 3|$ get?

a	b	a log2 – b log3	× a	× a ²
3	2	1.178e–01	0.3533	1.06
8	5	5.212e–02	0.4169	3.34
19	12	1.355e–02	0.2575	4.89
65	41	1.146e–02	0.7451	48.43
84	53	2.088e–03	0.1754	14.73
485	306	1.022e–03	0.4958	240.46
1054	665	4.365e–05	0.0460	48.50
24,727	15,601	1.819e–05	0.4499	11,124.62
50,508	31,867	7.265e–06	0.3669	18,533.29
125,743	79,335	3.665e–06	0.4608	57,943.85
176,251	111,202	3.600e–06	0.6345	111,839.07
301,994	190,537	6.449e–08	0.0195	5,881.89

(Every row is a convergent of $\log 3 / \log 2$ and reproduces from `tower_phase_spine.py` with its search ceiling raised from 4,000 to 400,000.)

The deficit shrinks by seven orders of magnitude. **Neither product follows it down.** Multiplied by a , the value wanders between 0.02 and 0.75 with no trend. Multiplied by a^2 , it climbs by five orders of magnitude and never comes near zero.

That is not a pattern we noticed. It is the shadow of **Baker’s theorem on linear forms in logarithms** (Baker 1966; effective constants, Matveev 2000), which gives a computable lower bound of the form $|a \log p - b \log q| > C \cdot \max(a,b)^{-\kappa}$, with κ **strictly greater than one**.

Sharper still, and specific to this exact form — a 2026-08-22 prior-art sweep found the operative bound this section should lean on. For rational combinations of $\log 2$ and $\log 3$, Padé-approximant methods beat the general machinery by orders of magnitude in the exponent: Rhin (1987) proves $|u_0 + u_1 \log 2 + u_2 \log 3| \geq H^{-\{7.616\}}$ for $H = \max(|u_1|, |u_2|)$ large, so the irrationality exponent of any nonzero rational combination of $\log 2$ and $\log 3$ is below 8.616 (Zudilin 2004, Thm. 3), refined to 7.6155 by Wu (2003). In plain words: the floor under the closure deficit is not just Baker’s qualitative promise backed by Matveev’s enormous general constant — for this paper’s two logarithms there is a small, explicit, published exponent, and the deficit provably cannot fall faster than a fixed power of the coefficient. The general Baker-Matveev route remains the fallback the moment a third prime enters the form.

That exponent is the whole point, and the table shows why it has to be there. The theorem bounds the deficit against a power of a *above* one, so it is the a^2 column that tracks the guarantee. The a column is a weaker empirical companion: it stays order one across everything we searched, but its floor is a fact about the continued fraction of $\log 3 / \log 2$ rather than a consequence of the theorem. We report both columns so the distinction is visible rather than assumed.




Figure 4. The floor, three ways. *Left:* the closure deficit $|a \log 2 - b \log 3|$ falling across the convergents. *Centre:* the same deficit multiplied by a , which stays order one and never reaches the zero line. *Right:* the diagnostic — that product plotted against $1/q$, where q is the next partial quotient of $\log 3 / \log 2$, showing what governs the centre panel. The guarantee the theorem supplies is on the a^2 column of the table above; the centre panel is the empirical companion.

So a nonzero closure deficit on one prime line cannot be driven arbitrarily close to zero by riding another. Zero is isolated, with a provable gap around it.

This is the only substrate in the whole program where the noise floor is proved, effective, and

uniform rather than measured. And it is not the inherited-property trap that catches the golden-angle arguments elsewhere in this work — the golden angle was my design choice, so anything special about it is something I put there. **Log 2 and log 3 are not a choice. They are what the primes are.**

THE LEFTOVER IS BENFORD, AND WE PREDICTED WRONG

Pick a base prime p and ask how many p -steps reach n in log space:

$$w = \log n / \log p = \text{whole wraps} + \text{leftover}$$

The leftover is zero exactly when n is a power of p . Everything else carries a defect it can never discharge — multiply by p all day and the wrap count goes up by one each time while the leftover does not move at all. It depends only on the part of n that has no p in it.

And the leftover **adds** under multiplication: $\{\log_3 25\} = 2 \times \{\log_3 5\}$ and $\{\log_3 35\} = \{\log_3 5\} + \{\log_3 7\}$, both to 0.00e+00.

The registered prediction failed, and failing was the informative outcome

We registered that the leftover would be **evenly spread**, and therefore carry no information. That was wrong, and the run said so:

range	deviation from EVEN	deviation from BENFORD	
2...10 ³	0.0451	0.0109	4.1× better
2...10 ⁴	0.0423	0.0153	2.8× better
2...10 ⁵	0.0430	0.0192	2.2× better

Deviation from even is **flat in N** — 0.0451, 0.0423, 0.0430. Even spreading would force it down. It doesn't go down and it can't.

The mistake was mine and it is a common one: Weyl's equidistribution theorem governs $\{n \cdot \alpha\}$, **not** $\{\log n\}$. Log grows too slowly.

What it actually follows is Benford's law in base p — the same law that governs leading digits in real-world data — cyclically rotated by the fractional part of $\log_p N$. At $N = 10^5$ the rotation is 0.4795 and the observed cliff lands exactly where that predicts.

That $\{\log n\}$ fails to equidistribute mod 1 is classical (Kuipers & Niederreiter 1974), and the dictionary between Benford's law and uniform distribution of the log mod 1 is Diaconis (1977); the rotated finite-N packaging measured here is ours, held as an observation on classical ground.

And this is not a footnote. Benford's law is precisely the distribution that is invariant under multiplying everything by a constant. We *separately* measured that the leftover adds under multiplication. Benford is the equilibrium distribution of exactly that operation. **The two results are one result:** the leftover lives on a circle carrying an additive group action, and the logarithmic measure is that action's invariant measure.

Had the leftover been evenly spread it would have carried nothing. Being Benford, it carries the multiplicative symmetry of the integers themselves.

THE OLDEST CLOSURE DEFECT ANYONE MEASURED IS MUSICAL

$\log 3 / \log 2 = 1.5849625007$. Its continued fraction is [1,1,1,2,2,3,1,5,2,23], and its convergents are the musical temperaments:

ratio	value	error	what it is
8/5	1.60000000	1.50e−02	5-tone
19/12	1.58333333	−1.63e−03	12-tone equal temperament
65/41	1.58536585	4.03e−04	41-tone
84/53	1.58490566	−5.68e−05	53-tone, Mercator's

$$3^{12}/2^{19} = 1.013643265 = \textbf{23.4600 cents.}$$
 The Pythagorean comma.

Twelve perfect fifths do not close into seven octaves. They miss, and *The Floor Is a Theorem* says they must — closing would require $\log 3 / \log 2$ to be rational.

Every tuning system in history is a decision about what to do with a defect that provably will not

go away. Spread it evenly across all twelve keys. Dump it into one interval nobody plays — the wolf fifth. Or add steps until it drops below hearing, which is the 53-tone system.

Distribute, localise, or refine. Those are the same three options an engineer has when a measured holonomy will not close, and they were enumerated two thousand years before anyone had the vocabulary. Jing Fang (78-37 BCE) reached the 53-tone convergent without logarithms. He had the twelve-fifths comma, extended the same calculation out to sixty fifths, and found that after fifty-three the new pitches were landing almost exactly on ones he had already computed. He put the leftover at 177147/176776 — six digits, by hand, two thousand years before anyone wrote down a logarithm. The same observation was made again in Europe by Nicolas Mercator (c. 1620-1687), which is why $3^{53}/2^{84}$ carries his name in the table above. (McClain, E. G. & Hung, M. S., *Chinese Cyclic Tunings in Late Antiquity*, **Ethnomusicology** 23(2), 1979, 205-224.)

A KNOT MEASURED ON BOTH SIDES

Knot determinants are integers, so a knot has a tower address. Under connected sum determinants multiply, so **addresses add in log space.**

Measured on PDB 8B1N (TrmD-Tm1570), 151 closures × 4 seeds, predictions registered before the run:

what	det	tower rung
first domain	3	log 3 = 1.0986
second domain	3	log 3 = 1.0986
whole chain	9	log 9 = 2.1972

log 3 + log 3 = log 9, residual **0.000e+00** — exact, because determinant multiplicativity is a theorem. **Both sides of the identity were measured on the same molecule**, not inferred from one side.

The contrast is the finding. The same molecule also carries writhe, which is also additive: 16.394 + 13.904 = 30.298 against a measured 30.470. Residual **+0.172, or 0.56%.**

channel	why it adds	leftover
determinant	it multiplies — a theorem	0.000e+00
writhe	it’s an integral; pieces sum	+0.172

Two additive laws on one object. The integer channel’s addition is exact and algebraic; the continuous channel’s is approximate and analytic, and its leftover measures the geometric coupling at the junction.

Fence: n = 1. One measurement. Weakest claim in the paper.

Limit, stated where found: det = 9 does not name the knot — 6₁ and 3₁#3₁ share it. The tower gives an address, never a resident.

RESONANCE AND MUSIC — AN OBSERVATION THAT CATCHES ITSELF

A climb is a log-ratio, and prime log-ratios are a musical scale. The same dial the companion paper describes — simple ratio, orbit closes, consonance; golden angle, never closes, dissonance — is in physics called **mode-locking**, or Arnold tongues.

Where it touches established work: KAM theory and Greene’s criterion, Arnold tongues in circle maps, phyllotaxis, orbital resonances, and the Hofstadter butterfly.

And here is the catch, which we ran on ourselves. The tower is *built* with the golden angle. So any golden-mean extremality it shows is **inherited, not discovered**. Four of those five collapse to that inheritance the moment a specialist looks. Only Hofstadter is independent, because both objects are continued-fraction-indexed for unrelated reasons.

Honest framing: one observation, not five discoveries. The golden angle places the tower on the same maximally-irrational point that organizes phyllotaxis, KAM, and Arnold tongues. A mathematician would see the inheritance instantly and should.

That contrast is worth stating against *The Floor Is a Theorem*: that section is **not** inherited, because the primes were not chosen.

THREADS LEFT DELIBERATELY IN THE OPEN — A HANDOFF

I do not usually return to a paper once it is finished, so anything I cannot close myself gets said out loud here instead of dying in a drawer. These are open invitations. Everything needed to pick them up — the construction, the code, the data — ships in this record's supplement under CC BY 4.0.

The diffraction pattern of this point set has never been computed, because the point set did not exist until this construction built it. The prime-rank golden-angle embedding — headings assigned by the rank of the prime among primes times 137.50776° , not by n itself — is, as far as a 2026-08-22 literature sweep could determine, a new deterministic aperiodic point set. There is a research program built precisely for such objects: the study of Vogel-spiral and aperiodic photonic arrays — structure factors, diffraction, hyperuniformity classes, wave localization — developed by Luca Dal Negro's group at Boston University and collaborators (Trevino, Cao & Dal Negro, *Nano Letters* 11, 2011, 2008–2016; Dal Negro, Wang & Pinheiro, *Crystals* 6, 2016, 161; Sgrignuoli, Wang, Pinheiro & Dal Negro, *Phys. Rev. B* 99, 2019, 104202; and, closest in spirit, Dal Negro, Henderson & Sgrignuoli, "Wave Transport and Localization in Prime Number Landscapes," *Frontiers in Physics* 9, 2021, 723456). The standard Vogel spiral indexes points by n ; this construction indexes headings by prime rank, so the azimuthal structure carries the arithmetic — including the rotated Benford phase measured above — and nobody knows what its diffraction pattern looks like. That question belongs to their instruments, not mine. If it shows the Vogel spiral's diffuse circular rings, that is an answer; if it shows arithmetic azimuthal structure, that is a better one. Consider this a warm handoff: the point set regenerates deterministically from the supplement's code, and I would be glad for it to become someone's afternoon.

The sideways constant wants an analyst. -0.32774696 at $s = 2$ has no located closed form; the natural attack — discrepancy control of the golden-angle partial sums via the Erdős-Turán inequality and the three-distance theorem (Sós 1958; Marklof & Strömbergsson, *Amer. Math. Monthly* 124, 2017) — should yield rigorous convergence rates and possibly an asymptotic as $s \rightarrow 1^+$. Whether it is irrational is open, and probably hard, and exactly the kind of hard that attracts the right person.

One rhyme, fenced. The tower and the Hofstadter butterfly are both continued-fraction-indexed objects in which the golden mean plays an extremal role (Avila & Jitomirskaya, *Ann. of Math.* 170, 2009, 303–342; Koch, arXiv:1907.06804). A 2026-08-22 sweep found no published bridge between a multiplicative number theory construction and the butterfly, and none is claimed here — this sentence exists so that a reader who feels the same pull knows the pull has been felt, and that it is, for now, only a pull.

STANDING BESIDE THE PRIOR WORK

Everything load-bearing under this construction was discovered and named by someone else, and we found several of them *after* rederiving them.

Classical, cited, not claimed. The integers under multiplication as a free abelian group on the primes. The Liouville function and its Dirichlet series $\zeta(2s)/\zeta(s)$ — textbook (Apostol; OEIS A008836). The arithmetic taxicab metric $\Omega(\text{lcm/gcd})$ — **published** (D. Dominici, arXiv 0906.0632). Effective bounds on linear forms in logarithms — Baker 1966, Matveev 2000; the sharp two-log measure — Rhin 1987, Wu 2003, Zudilin 2004. Benford's law — Newcomb 1881, Benford 1938, Diaconis 1977, Kuipers-Niederreiter 1974. The divides-relation graph, invariant by invariant — Cha, DuCasse & Quintas 2014. Maximal-chain counts as linear extensions — Stanley, *Enumerative Combinatorics* I. Equidistribution — Weyl 1916, cited here for what it does *not* cover. The factorization-torus — Lulek & Olchawa 1996. Continued-fraction convergents as temperaments — the entire tuning literature.

Asserted only as "not located." The chord $H(n)$ as a two-dimensional character sum whose frequencies are the golden angle indexed by prime rank. Constant -0.32774696 at $s = 2$.

And a caveat that must ship with that. The literature neighbourhood around "geometric prime factorization" contains a great deal of non-peer-reviewed material, so a not-located result from a general search engine is weaker evidence here than it would be elsewhere. The fix is a reviewed-literature index rather than a search engine. zbMATH Open indexes reviewed mathematics back to 1868 and has been free to search since 2021, which makes it exactly the filter this neighbourhood needs. **That pass has not been run.** Until it has, no novelty language beyond *not located*.

THINGS THAT CAN LOSE

The vertical prediction. The eight sideways regularities reproduce identically under any vertical; the octave law does not. *Dies if* any sideways regularity changes when the vertical changes.

The Baker floor. For distinct primes p, q there are effective constants C and $\kappa > 1$ with $|a \log p - b \log q| > C \cdot \max(a, b)^{-\kappa}$, so zero is isolated with a computable gap around it. *Dies if* any pair admits a nonzero deficit with no effective lower bound. The weaker companion observation — that deficit $\times a$ stays order one — is empirical, holds everywhere we searched out to $a \approx 3 \times 10^5$, and would fail at a sufficiently large partial quotient without touching the result above.

Benford everywhere. The leftover follows Benford base p , rotated by $\text{frac}(\log_p N)$, for every prime p . *Dies if* any prime's leftover fits even spreading better.

The $n-1$ bracket. Holds for every prime except 5. *Dies at* the first other prime whose reach escapes its window.

HOW TO BREAK THIS

Every number here regenerates from `tower_walk_cost.py` and `tower_phase_spine.py` — pure standard library, no dependencies, they write their own log files. Run them. If a number disagrees with this paper, that is the paper's problem and we want to know. Breaks get logged in the open with attribution.

The interesting attacks are not arithmetic. They are: *is the golden angle the right connection or an arbitrary one? Does grazing-equals-divisibility reveal structure in the integers, or only in my embedding? Is the closure deficit measuring number theory, or measuring my coordinate system?* The math does not settle those. It only guarantees the argument can happen on solid ground.

DEPOSIT METADATA

Author Schulz, Matthew · Independent Researcher · Belmont, CA **ORCID** 0009-0000-0956-5743 · **Contact** mattschulz97@protonmail.com **License** CC BY 4.0 · **Type** Preprint · **Date** 2026-08-22 (v3.4; v3.3 deposited 2026-08-01)

Related deposits: *Taking a Loop Walk Through the Integer Forest* 10.5281/zenodo.21449569 · *Walking the Loop: A Quick Guide to Holonomy* 10.5281/zenodo.21448986 · *Reading the Sky Through Loops* 10.5281/zenodo.21331651

Keywords multiplicative number theory; prime factorization; golden angle; Liouville function; Dirichlet series; Laplace transform; Baker's theorem; linear forms in logarithms; Benford's law; musical temperament; coordinate-free; closure deficit