

The Magnetic Forest: a two-channel instrument for the integer orchestra, the gas it turns out to be, and the price of a coincidence

Matthew "Bear" Schulz

Independent researcher · Belmont, California

ORCID: 0009-0000-0956-5743

Licence: CC BY 4.0

Preprint · v2 · 26 August 2026

Every figure and every number in this paper regenerates from a receipt shipped in the accompanying crate. Where a claim is not measured, it carries a tag saying so.

Keywords: magnetic graph Laplacian, lattice gauge theory, primon gas, smooth numbers, Baker's theorem, continued fractions, Sparrow criterion, resolution limits, persistent current

Abstract

We build a two-channel instrument that reads the multiplicative structure of the integers the way an interferometer reads a light source, discover that the object it reads is a gas, and use both facts to measure what a near-coincidence between two prime towers costs to *see*.

The instrument's channels are a magnetic graph Laplacian on the tower network of smooth numbers and an open-beam interference screen indexed by the same towers. A design finding governs both: a character on the tower graph is pure gauge, so the spectrum must not move when characters are swept. That non-result is the instrument's built-in correctness receipt, measured at $2.04\text{e-}14$. Structure enters only as flux through the graph's faces, and it does — the spectrum moves by 0.725, is periodic in the flux, and splits into bands at rational flux.

On the open channel each prime tower is a comb of resonances, and the near-coincidence between two towers — the Pythagorean comma, 2^{19} against 3^{12} — appears as two teeth 0.111810 apart. A 142-node aperture cannot separate them. Chasing that failure produced the rest of the paper.

Weighting a tower node by n to the minus sigma is a Boltzmann factor with energy $\log n$, so the forest is the two-species primon gas; its mean energy turns out to be Bose-Einstein occupation per species, and its partition function is exactly solvable. It has no Hagedorn wall of its own — a two-species gas cannot — while the full integer gas does, appearing here as a heat-capacity knee that climbs through 5.99, 9.38, 13.62 and 18.73 as the patch grows through four decades. Crucially, that same weight is an *apodization* of the aperture, which makes temperature a resolution parameter.

We then derive a closed form for the beam near any near-coincidence costing three exponentials rather than one term per lattice node, validate it against the honest node sum wherever that still runs, and reach a coincidence whose aperture would require summing $1.3\text{e}14$ nodes. In scaled coordinates every near-coincidence in every prime pair collapses onto one universal curve, and the barrel at which one peak becomes two — Sparrow's 1916 criterion, which has no free parameter — is the first root of one scalar function. We solve it: **1.20871040964352282**. A census of 993 prime pairs and 14,098 rungs measures 13,978 of them; of the 8,295 deep rungs, not one lands above the solved value.

But that number is the infinite-temperature limit of a curve, and the curve ends. Beyond a critical damping no barrel of any length separates the two teeth, and **the critical value is exactly π** — proved by a one-line factorization and confirmed to $1.4\text{e-}07$. So a coincidence is visible only if the gas is hotter than the barrel it demands divided by π , and every finite temperature cuts the ladder of coincidences at a finite depth. The census's most charming find makes the point without apparatus: $13^7 = 62,748,517$ and $89^4 = 62,742,241$ differ by one part in ten thousand, findable with a pocket calculator, and resolving them would take a telescope 115,100 logarithmic units long.

We claim the instrument, the census, and the number. The criterion is Sparrow's, the floor is Baker's, the bargain distribution is Gauss and Kuzmin's, and the gas is Julia's. We end with four ways to break the result.

0. Where this comes from

This paper is the fourth panel of a sequence, and the sequence has a shape worth stating because it explains the method rather than decorating it.

The *Log and Anti-Log (Linear) Towers of the Integer Forest* established the object: the integers arranged by their multiplicative structure into towers, walked rather than summed. *Loop Walking Cascadia* took the same walking method outdoors, wrapping a real GPS network in a lattice, treating the stations as a fiber bundle and reading what the loops — not the stations — had to say. *A Coincidence Has No Single Price* established that a coincidence's improbability is not a single number: what it costs depends on the register you compute it in, a lesson learned by revisiting a dark-energy exponent argument and finding its price changed under re-registration. And *The Crossover Atlas* walked known walls in languages they were not built in, including the chapter that identified the Pythagorean comma as a row of the Baker-floor table — a chapter that deliberately downgraded its own verdict to settled-and-classical after finding Barbour had been there in 1948.

The present paper is what happens when the Cascadia move is turned back on the Integer Forest: having wrapped a real network in a lattice and listened to its loops, wrap the towers the same way and listen to theirs.

One correction to that impulse, because it would otherwise mislead. Cascadia's lattice was made of triangles. The forest's is not. The smallest closed loop here runs from n to $2n$ to $6n$ to $3n$ and back, and that is a four-sided face; this patch has 110 of them, and they are squares, not triangles. What carries over from Cascadia is not the triangle but the principle underneath it — that the instrument is read on the *faces*, not the nodes, because a node value is exactly the part a gauge choice can erase. The shape of the face is set by the arithmetic, and two prime directions give a quadrilateral.

The other continuity is the price. *A Coincidence Has No Single Price* argued that a coincidence's cost depends on how you register it. Part III of this paper makes that literal: the price acquires units — a length of telescope — and turns out to depend on one register we had not considered, the temperature of the gas the coincidence lives in.

Concept DOIs for the works named above: the *Integer Forest*, [10.5281/zenodo.21449569](https://doi.org/10.5281/zenodo.21449569); *The Tower Is Coordinate-Free*, [10.5281/zenodo.21798310](https://doi.org/10.5281/zenodo.21798310); *Loop Walking Cascadia*, [10.5281/zenodo.22076451](https://doi.org/10.5281/zenodo.22076451); *A Coincidence Has No Single Price*, [10.5281/zenodo.22053824](https://doi.org/10.5281/zenodo.22053824); *The Crossover Atlas*, [10.5281/zenodo.22064125](https://doi.org/10.5281/zenodo.22064125).

PART I — THE INSTRUMENT

1. The forest, and why a character cannot matter

Take the integers of the form $2^a \cdot 3^b$ below a million — 142 of them — as the nodes of a graph, and join n to $2n$ and to $3n$. This is the tower network: two ladders leaning on each other, every node reachable by some number of doublings and triplings. On its edges we hang phases, making the adjacency matrix Hermitian, and we read the eigenvalues. That object is a magnetic graph Laplacian, and the phases are a connection on a bundle over the graph.

The first thing to establish is what the phases *cannot* do. If each prime carries its own angle — the times-2 edge always stamped with one angle, the times-3 edge with another — then going around the smallest closed loop, from n to $2n$ to $6n$ to $3n$ and back to n , picks up each angle once forward and once backward, which is nothing. The phases cancel exactly around every commuting face. In gauge language the character is pure gauge: re-phasing the nodes absorbs the whole dressing, and the spectrum cannot know it happened.

This is easy to mistake for a disappointment. It is the opposite: it is the instrument's built-in correctness test, the one measurement whose *failure* would mean the apparatus is wrong. Sweeping the characters across their full range moves the spectrum by at most **2.04e-14**, and dressing every node with an independent random phase on top moves it by at most **2.00e-14** — machine zero in both cases, against a declared tolerance of $1e-10$. The bundle is flat under dressing, measured rather than assumed.

2. Flux, and the butterfly

Structure enters when the phase around a face does not cancel. We make the times-2 edge phase depend on where the node sits — specifically on its power of three — so that every one of the 110 faces carries the same flux. Now the loop integral is not zero, and the spectrum has something to respond to.

It responds. At unit flux the eigenvalues move by up to **0.725**, against the $1e-14$ they moved under any amount of character dressing, and the motion depends only on the flux and not on which gauge realizes it. The physics lives in the flux, never in the dressing.

Sweeping the flux through a full turn and stacking the spectra produces Figure 1, the forest's butterfly. Bands split where the flux is a rational multiple of a turn and smear where it is not. This is a kinship with Hofstadter's spectrum for electrons on a lattice in a magnetic field, and we report it as a kinship: the qualitative signature is there, the forest is not a square lattice in the physical sense, and we make no claim of identity.




Figure 1. The forest's butterfly: eigenvalues of the 142-node tower graph against flux per face, swept over one full turn. Band structure appears at rational flux. [kinship, fenced]

3. The open beam: combs, the comma, and the silences

The same object has an optical reading. Index a screen by tower address — the cell at (a, b) standing for the integer $2^a \cdot 3^b$ — and give each cell the phase t times $\log n$. The far field is the Fraunhofer transform of that screen, and the on-axis amplitude is the coherent sum of $\exp(i \cdot t \cdot \log n)$ over every node. In words: the screen encodes the whole orchestra of towers, and t is the tuning dial. Turning the dial asks which chord the integers are playing at that frequency.

The flatness result reappears in open-beam form: changing the characters translates the far-field pattern without altering its structure, exact to **9.88e-14**, and to $7.22e-16$ when the shift lands on a whole sampling bin. A character moves where the picture sits, not what is in it.

Turning the dial reveals that each tower is a comb. Tower 2 chimes whenever t times $\log 2$ is a whole number of turns — every 9.0647 in dial units — and tower 3 every 5.7192. The two combs mostly ignore each other, and then occasionally nearly agree. The most famous near-agreement is the Pythagorean comma: twelve stories up tower 3

lands almost exactly nineteen up tower 2, because $3^{12} = 531,441$ and $2^{19} = 524,288$ sit within a whisker of each other logarithmically. On the dial this is tower 2's twelfth tooth at $t = 108.7766$ and tower 3's nineteenth at 108.6648 : two lines **0.111810** apart, that number being exactly 2π times the logarithmic near-miss divided by the product of the two logarithms.

The 142-node screen cannot separate them. Its resolving power — a full turn divided by the range of logarithms it spans — is 0.455, four times coarser than the gap. The teeth fuse into one peak at $t = 108.7206$, and that peak stands at 0.946 of full coherence, meaning almost the entire patch is marching in step. Nothing is synchronizing; that one dial setting simply makes the phase step in the doubling direction and in the tripling direction both land near a whole turn at once, so every node's phase falls inside a narrow arc. It is a coincidence of one dial setting, not a lock.




Figure 2. The tuning dial: coherent amplitude against t across the tower combs, with each string's harmonics marked and the comma's fused peak at $t \approx 108.72$.

Between the peaks are the silences, and they are informative in an unexpected way. With only two towers on the screen the beam never goes quiet: its minimum over the swept range is **0.371**, comfortably above zero. Silence requires everyone. When the full orchestra of integers is on the screen the beam does dip near the first few nontrivial zeros of the Riemann zeta function — unsurprising, since the beam is a truncated zeta sum along a vertical line, and we tag it as a finite-aperture kinship rather than a result. A truncated sum is not the function.




Figure 3. The silences: minimum amplitude across the sweep for the two-string screen and for the full-orchestra beam. Two strings never cancel; cancellation is a collective act. [kinship, fenced]

Finally, drawing the same flux on the screen — a phase proportional to the product of the two tower addresses — reorganizes the far field, moving it by 0.957 of the peak at a flux of 1.3, and returns it after a full turn. The two channels agree: dressing does nothing, flux does everything, and the response is periodic.

4. The comma drawn three ways

Before leaving the instrument we draw the comma itself three ways, because the rest of the paper is about it. We claim none of these; the identification of the comma with a row of the Baker-floor table was walked in the Crossover Atlas and downgraded there to settled-and-classical on finding Barbour had done it in 1948.

As a pair of rulers, the two towers are vernier scales with incommensurable spacings: they approach, nearly touch at the twelfth story, and never touch, because unique factorization forbids $2^a = 3^b$ for positive exponents. The near-miss measures **1.355103e-2** logarithmically.

As a dial, the comma is a precession. Each story up tower 3 lands at a fractional position on tower 2's octave circle, and those landings drift. After twelve stories the drift is 23.46 cents — the Pythagorean comma of music theory, reached from the arithmetic rather than the ear. After fifty-three stories it is 3.615 cents, Mercator's comma, the reason 53-tone temperament is nearly perfect.

As a ladder, the record near-returns occur at stories **1, 2, 5, 12, 41, 53, 306** and onward, exactly the convergent denominators of log base 2 of 3. The ladder never bottoms out.




Figure 4. The comma drawn three ways: vernier rulers that approach and never touch; the precession dial through the Pythagorean and Mercator commas; the ladder of record near-returns.

That last picture is the hinge. If there is always a better coincidence further up, and the 142-node screen could not resolve even the twelfth-story one, the question is not whether these coincidences exist. It is what they cost to see. Answering that turned out to require knowing what the forest is thermodynamically, which is Part II.

PART II — THE GAS

5. The forest is a two-mode Bose gas

Weighting a node by n to the minus sigma is a Boltzmann factor, because the energy of the state n is $\log n$: each prime is a particle species whose quantum costs $\log p$, and the node $2^a \cdot 3^b$ is a state holding a quanta of the 2-particle and b of the 3-particle. This is the primon gas, and every beam sum in this program has been a partition function wearing radio clothes.

One convention must be stated once, because it is easy to get backwards and we did: sigma is the *inverse* temperature. Large sigma is cold. Sigma equal to zero — the undamped beam used throughout Part I — is *infinite* temperature.

Because the two species are independent the partition function factorizes into two Euler factors and is exactly solvable, so free energy, mean energy and heat capacity are all closed-form and the finite patch has something real to fail against. It does not fail: the patch sum reproduces the closed form to **5.55e-16**, and the heat capacity computed from the closed form agrees with the heat capacity differenced from the energy to 1.25e-07.

The mean energy is worth stating separately. Each species contributes its log times an occupation number, and that occupation number is exactly the Bose-Einstein form. The forest is a two-mode Bose gas, and its heat capacity tends to exactly 2 at infinite temperature — one unit per species.

6. The wall that is not ours

The primon gas is famous for a maximum temperature, and we expected to measure it here. We cannot, and the reason is a scope correction worth stating rather than hiding.

A Hagedorn wall exists when the density of states grows exponentially with energy, because an exponential density fought against a Boltzmann factor is exactly what makes a partition function refuse to converge. The full integer gas has that: there are about e^E integers below e^E . Our two-tower patch does not: the count of 3-smooth numbers below e^x grows like x squared, so the density grows only linearly in energy and the partition function converges at every positive temperature. **A two-species gas cannot have a Hagedorn wall. You need all the primes.** Measured, the two-species heat capacity peaks at a bounded 1.9996 and diverges nowhere.

So we measure the wall on the full integer patch instead — a different aperture from the forest, declared rather than smuggled. Summing every integer up to N and differencing, the heat-capacity knee climbs through **5.99, 9.38, 13.62 and 18.73** as N runs from a thousand to a million, tracking the square of $\log N$ over twelve, with its peak walking

inward from sigma 1.30 to 1.09 — toward sigma equal to one, the Hagedorn temperature, above which the full gas has no equilibrium at all.

Which places this program in an odd and honest spot. The undamped beam used throughout Part I sits at infinite temperature, far above that wall. It is finite only because our patch is truncated.




Figure 5. The gas: the exactly solvable two-species backbone with its bounded heat capacity (no wall); the full integer gas's knee sharpening toward the Hagedorn temperature as the patch grows; the splitting barrel as a function of temperature, ending at exactly π ; and the persistent current at three temperatures.

7. Temperature is apodization

Here is the fact that makes the gas load-bearing rather than decorative.

Written in the aperture's own coordinates, the Boltzmann weight is an exponential taper across the barrel: the far end of the instrument is dimmed by e to the minus sigma times distance. That is an apodization window, and apodization is a *resolution* question, not a thermodynamic one. Cooling the gas shortens the effective aperture. Temperature is therefore a parameter of Part III, and Part III cannot state its constant without naming which temperature it belongs to.

The instrument survives the change untouched. Cooling sends both tower angles complex, but the combination that defines the patch hypotenuse's lean — one angle minus the other times the ratio of the logs — has its imaginary parts cancel identically, so the lean stays exactly real. Checked rather than assumed: **0.00e+00**. The sum-free beam of Part III therefore works at any temperature with no new machinery, and it tracks the honest node sum to 6.1e-04.

8. The invariant riding on the thermodynamics

There is a trivial version of "the flux invariant survives temperature" which is not worth printing: the spectrum is already periodic in the flux, so any function of it is periodic and a thermal partition function being periodic is a corollary rather than a finding.

The object with content is the persistent current, minus the derivative of free energy with respect to flux. Measured on the 142-node forest as the spectral temperature rises through 0.5, 2 and 8, the current's amplitude falls through **0.289, 0.025 and 0.00049** and its harmonic content sharpens, while its period holds to **2.09e-14** throughout. The thermodynamics changes how loudly the invariant speaks; it never changes what it says. The resemblance to persistent currents in mesoscopic rings is a kinship and is fenced as one.

PART III — THE PRICE OF A COINCIDENCE

9. The aperture law, and where it comes from

Two spectral lines separated by a small distance need an aperture comparable to a full turn divided by that separation before an instrument can tell them apart. Applied here, the comma's 0.111810 gap suggests a logarithmic baseline of $2\pi/0.111810 = 56.19$. That is the right *scale* and not the right answer: measured against the honest node sum the teeth are still one peak at 56.19, and part at 67.0.

Where the correct number comes from falls out of the algebra rather than being imposed on it. Near a near-coincidence every node's phase collapses onto two tiny angles — how far the dial sits from tower 2's tooth, and how far from tower 3's — because the enormous integers involved are whole turns and whole turns are invisible. Summing over the patch then closes twice. Down each column the available stories of the second tower form an evenly spaced comb of phases, which has a closed form. Across the columns the patch's hypotenuse leans by a fixed amount independent of the dial, so what is left is a second comb. The result is a difference of two comb-sums divided by the tooth spacing: three exponentials, no sum at all.

Hiding in that expression is the aperture law. Multiply the hypotenuse's lean by the number of columns and you get exactly 2π times the barrel measured in units of the barrel the coincidence demands. So the second comb falls on its own first null when the barrel equals the naive prediction. The classical criterion is not assumed by this instrument; it drops out of it. And because the split happens a little past that null, the true threshold necessarily overshoots.

10. Removing the wall

The closed form is a claim, so we validate it against the node sum at a shallow rung where the approximation is at its worst and at the deepest rung whose columns can still be counted. The worst disagreement in the measured splitting barrel is $1.34\text{e-}2$ shallow and **5.81e-04** deep. The approximation is weakest exactly where we do not need it and strongest exactly where nothing else can go.

That lets us look at a coincidence no sum can reach. The (2,3) forest's deepest bargain within computational reach is 2^{301994} against 3^{190537} : after 190,537 stories the towers lean within $6.45\text{e-}8$ of each other, and unique factorization still forbids the touch. Its teeth sit on the dial at $t = 1,727,164.6087$ and are **5.32e-07** apart — agreeing in their first twelve significant figures and parting in the thirteenth. It demands a barrel of $1.18\text{e}7$ and splits at $1.43\text{e}7$, a patch standing for about **1.3e14 nodes**. Summing that is impossible on any machine; the closed form reads it in three exponentials.




Figure 6. Before and after: the comma fused at 142 rotors, resolved at a longer barrel, the transition as the barrel grows, and the ladder of barrels each rung demands.

11. The constant, which is a curve

Measure position in units of the gap and barrel in units of the barrel demanded, and the beam stops depending on *which* coincidence it is reading. Every near-kiss, in every prime pair, becomes the same curve, the distance from that limit shape falling from $1.5 \cdot 10^{-1}$ at the shallowest rung to $4.8 \cdot 10^{-9}$ at the deepest.

On a universal curve the splitting barrel is the root of one scalar function — but which root depends on which question is asked, and we got this wrong first. Our initial test asked whether two peaks had appeared *and landed within fifteen percent of a gap of their true positions*. That fifteen percent was a choice, and the answer moved with it: 1.219 at a thirty percent tolerance, 1.360 at one percent. Such a number is not a constant; it is the price of an accuracy somebody selected.

The question with no free parameter is older and is not ours. Sparrow, in 1916, held two lines resolved when the intensity at the midpoint between them first shows a minimum — the moment one peak becomes two. Expanding the modulus makes our profile a real analytic function of position, with no absolute value, no square root and no removable singularity, so Sparrow's condition becomes simply that the curvature at the midpoint passes through zero. Solving that rather than sampling it gives

1.20871040964352282

with a residual of $2.4 \cdot 10^{-46}$ and no drift between working precisions. Reading the same constant off a sampled grid is good only to about seven figures, a caution worth recording: we quoted more digits than we had earned before we solved it.

Which definition is used matters more here than it usually does, because this pupil is complex — the second term carries a phase — and the standard statement of the Sparrow limit is framed for real ones. We therefore name it rather than leave it implied: the Sparrow point is taken as the zero-crossing of the curvature at the midpoint of the **modulus-squared** profile, which expanding the modulus renders real analytic. That is what is solved above.

And by Part II this number now needs an address. It is the value at sigma equal to zero — the **infinite-temperature** limit of a curve. Cooling the gas dims the far end of the barrel, shrinking the effective aperture, so the splitting point rises: measured against the damping across the demanded barrel, it climbs through 1.2110, 1.2183, 1.2287, 1.2525, 1.2661 and on.

12. The census: 993 patches

Every prime pair is its own forest with its own ladder, floor, and required telescope. We surveyed 993 of them, recovering each ladder from raw logarithms — climb the tower, keep every story that lands closer to home than all before it — and checking it against the continued-fraction answer key. Across 993 patches, **no rung was missed**: the record-holders are exactly the convergent denominators, found by scanning rather than by asking.

That yields 14,098 rungs, of which 13,978 were split and priced. Over the **8,295 deep rungs**, where the finite-size effect has died, the measured Sparrow point runs from 1.208210 to 1.208710 and **not one lands above the solved value**. Every patch approaches the same ceiling from below and none crosses it — what a finite-size effect must look like and what a coincidence would not.




Figure 7. The 55-rung seen; the closed form against the exact node sum; the collapse of six rungs spanning six decades onto one curve; and the constant holding all the way down.




Figure 8. The census: each patch's ladder of bargains and ladder of telescopes; the distribution of bargains against the law of chance; and 13,978 measured rungs across 993 patches against the solved constant.

12.1 The bargains are not miracles

Rungs differ enormously in quality. We score each by how much better its near-miss is than a rung of its depth generically manages — a bargain factor, essentially the next partial quotient of the log ratio. The census champion is $79^{\wedge}75439$ against $157^{\wedge}65192$, leaning within $7.26\text{e-}9$ with a bargain factor of 9,226 and a telescope of three billion.

It would be easy to make too much of that. Gauss and Kuzmin settled what the partial quotients of a random irrational look like, and across all 14,098 rungs the census matches that law to a worst deviation of **0.0062**. No pair is blessed. The correct statement is not that some coincidences are miraculous but that coincidence has a *price distribution* whose tail is populated exactly as chance requires.

12.2 The one you can find with a calculator

The second-best bargain needs no apparatus at all. Thirteen to the seventh is 62,748,517; eighty-nine to the fourth is 62,742,241. They differ by 6,276 — one part in ten thousand — with a bargain factor of 6,411. Anyone with a calculator could stumble on it.

Its telescope is 115,100 logarithmic units. The most *findable* coincidence in the census is hopelessly beyond any instrument that could resolve it. You can notice a coincidence and still be unable to see it.

12.3 The clones, and why the census is complete

A worry the census does not obviously answer: every prime owns a copy of this entire forest. Multiply every integer by a prime P and nothing is distorted — the family of P is the integers wearing P as a badge, so the pattern of which tower each arrival lands in is just the factorisation of 1, 2, 3, 4, and plays out identically for every prime that exists. Might each of those countless families hide near-coincidences of its own that a survey of prime pairs never looked at?

It cannot, and the reason is one line. Because the logarithm of P times n is the logarithm of P plus the logarithm of n , the whole family is the origin's pattern slid along by a constant. Every near-coincidence inside a family is a difference of two of its members' logarithms, and the constant cancels from both. The clone's teeth therefore sit exactly where

the origin's do, with exactly the same Baker floor. Translation produces no new coincidences and destroys none, so a census over prime **pairs** is exhaustive over families. [exact]

The same shift says something about the instrument that is worth stating beside Part I's flatness receipt. Since every node of a clone picks up the same extra angle, the clone's beam is the original beam turned by a single global phase — measured here at $3.9\text{e-}11$ against the exact prediction, with the modulus untouched to $1\text{e-}12$. **Multiplying by a prime is a gauge transformation of the open beam.** The beam cannot see which prime's family it is looking at. What it can see is the barrel: inside a fixed forest, P's family reaches only as far as the cutoff divided by P, so a clone is shorter by the logarithm of P — measured at a cutoff of $1\text{e}18$ as 41.45 for the origin against 32.44, 27.35 and 17.22 for the families of 7,919, 1,299,709 and 32,416,190,071, with resolving powers degrading from 0.152 to 0.365 accordingly. In this instrument, being a different prime's family is not a difference in what you play. It is a difference in how long you get to play it.

None of this is new. In Dirichlet-series language it is the shift property, and structurally it is the free commutative monoid of the integers under multiplication acting on itself by translation. It appears here because its consequence — that the pair census is complete — is a question this program would otherwise have left open. What it does **not** settle is whether two different families can align with each other; that is the pair census again with both primes large, and it is not answered here. [open]

12.4 The test at giant primes

Section 12.3 leaves one question open: whether two *different* families can align with each other. That is the pair census again with both primes large, and it is the sharpest available test of the constant, because giant primes are where the claim is least exercised and hardest to check. We ran it as a falsification.

The check has a boundary that must be declared before any measurement rather than discovered after. The independent cross-check — the finite-aperture form — needs the exact column count of the aperture, and a double-precision integer stops counting exactly past 2^{53} , about $9.0\text{e}15$. Beyond that the universal shape can still be evaluated, since it never mentions the column count and grows more exact with depth, but the cross-check is gone. That is an honest asymmetry: the giant-prime regime is where our best tool is most trustworthy and our check least available. In the event, no rung of the four pairs tested crossed the boundary.

Twenty-seven rungs across four giant pairs were measurable. They divide cleanly. Thirteen rungs with fewer than ten thousand columns show the finite-size approach this paper already documents, climbing from 1.058 at an aperture of eight columns through 1.148, 1.204 and 1.207 as the aperture grows, always from below. The remaining fourteen, where the asymptotic claim actually claims something, agree with the predicted constant to a worst deviation of **6.46e-05** — and the deepest of them carries **1.25e10 columns**, an aperture of ten billion terms read in three exponentials. The prediction survived a test built to kill it.

We record one methodological stumble, because it nearly became a false result. Our first version of this test compared every measurable rung against the asymptotic value without regard to aperture, so the shallowest — an eight-column rung — reported a deviation of 0.151 and the ledger printed that universality was in trouble. It was not; the test was wrong, comparing a finite-size measurement against an infinite-aperture prediction. The split above is the correction.




Figure 9. The giant-prime test: where the independent cross-check would run out against the barrel each rung demands, and the temperature each rung requires.

The pairs also settle which giant primes are worth looking at, and the answer is counter-intuitive. Two primes that are close in *value* are useless: 1,000,000,007 against 1,000,000,009 have logarithms that nearly coincide, so their very first rung demands a barrel of $2.15e11$ and a gas at temperature $6.8e10$. Two primes whose *logarithms* are far apart are cheap: 1,000,003 against 32,416,190,071 has a first rung demanding a barrel of 32.19, visible at temperature 10.2. Four orders of magnitude in the primes; ten orders in the bill.

One last distinction, because the clone picture can be pushed past where it holds. Following a prime's multiplicative children — P , $2P$, $3P$, $4P$ — the walk through the towers is the factorisation of 1, 2, 3, 4, identical for every prime that exists. But the set of integers a prime owns as *their largest factor* is P times the P -smooth numbers, and that forest gains a tower for every prime below P : two towers at $P = 3$, twenty-five at 97, a thousand at 7,919, a hundred thousand at 1,299,709. One family is universal; the other grows without bound. Both are true of the same prime, and only the first is a fixed walk. [exact]

13. The price of knowing where

Detection has a second tier, found by accident of our own error. Knowing there are two teeth costs one number. Knowing *where* they are is a further purchase with a price curve: locating a tooth to within a tenth of a gap costs 1.27 barrels, to a twentieth 2.20, to a hundredth 5.15, to a two-hundredth 7.15 — approximately 0.5 divided by the square root of the accuracy demanded, diverging, so perfect localization costs an infinite telescope. The approach is not monotone either; the located tooth wobbles as the barrel grows, so the honest question is where the error *stays* below a tolerance. The fifteen percent tolerance that gave us our wrong constant was quietly buying an accuracy, and 1.2476 was its price.

14. The thermal horizon

The price story has a boundary, and it is exact.

Push the barrel to infinity at fixed temperature and the interference term dies, being suppressed exponentially in the damping across the actual barrel. What remains is a shape that is one over the square root of the product of two quadratics, one centred on each tooth, each with a width set by the temperature. Its derivative factors exactly into the midpoint root times a quadratic, and that quadratic's roots — the two side maxima, which are the two resolved teeth — are real if and only if a certain discriminant is non-negative. Written in terms of the damping measured across the barrel the coincidence demands, the condition is that this damping not exceed π .

Beyond π , no barrel of any length separates the two teeth. Bisecting the profile numerically agrees to $1.4e-07$.

Two appearances of π , and they are not the same π . The aperture identity of Section 9 carries a 2π because a full turn is 2π — pure convention, present only because angles are measured in radians. The horizon's π is *inherited from that same convention*, not independent of it: the merge condition is that the half-width not exceed half the separation, $a \leq 1/2$, and since the damping is measured in radians the threshold reads $2\pi \cdot (1/2) = \pi$. **In units of turns the**

horizon sits at exactly one half. We state this plainly because "a wall at exactly π " invites a reader to hear something transcendental, and there is nothing transcendental here. The nontrivial number is one half, and the algebra producing it is elementary.

Read plainly: a coincidence is resolvable only if the gas is hotter than the barrel it demands divided by π . Since the ladder's demanded barrels grow without bound, **every finite temperature cuts the ladder at a finite depth.** At temperature 1 only one of the first six (2,3) rungs survives; at 10, two; at 100, four; at 1000, all six. Deeper coincidences still exist, still obey Baker's floor, and are permanently invisible at that temperature.

We must be careful what kind of wall this is, because it is easy to overstate and we did at first. The horizon is a limit on *this criterion* — on resolving by counting peaks, which is what Sparrow's criterion means and what an imaging instrument does. It is not an information-theoretic impossibility. With the lineshape known analytically and no noise, a model-fitting estimator can recover both tooth positions from a single unresolved bump at any damping whatever; with noise it faces a different and softer bound of the Cramér-Rao kind. So the honest sentence is that beyond π the two teeth are invisible **to peak-counting**, which is the sense in which every classical resolution limit is a limit. What a fitted model could extract from the same data is a separate question this paper does not answer. [open]

One juxtaposition we deliberately do not multiply: the full integer gas cannot be heated past temperature 1 at all. If one insisted on that gas as the bath, only coincidences demanding a barrel under π would ever be resolvable — the first two rungs and nothing else. But our patch is two species and has no such cap, so the two statements belong to different apertures and we leave them side by side. [open]

14.1 The array question, and which currency the horizon charges in

Anyone who has built a phased array asks the same question of our filled aperture: resolution is set by the *longest* baseline, so why carry every node?

The question is well posed here, because the aperture is grotesquely redundant. What a correlator sees is not elements but baselines — here, differences of $\log n$ — and counting the distinct ones, the 142-node patch carries 19,881 ordered pairs across only 737 distinct baselines, a redundancy of 27. Elements grow with the square of the barrel while distinct baselines grow only linearly, so the redundancy climbs to 52 at a barrel of 20, 110 at 30 and **418 at 60**. The filled aperture is buying the same baselines over and over.

The budget array is immediate: keep the two bare towers and discard every mixed node. Same longest baseline, element count linear rather than quadratic — 1,753 elements instead of 365,306 at a barrel of 745, and at the deepest rung 3.3e8 instead of 1.4e14, **four hundred thousand times fewer**.

It does not come free, and this paper already met the failure once: the two-tower array at the naive barrel reports a pair 37 percent too wide, which is sidelobe structure wearing a mask. Measured across barrel lengths it converges to the true separation slowly and not monotonically — 14.7 percent wrong at 1.21 barrels, still 4 percent at four, 0.7 percent at eight, right at twelve. A thinned array needs roughly ten times the barrel to say what the filled one says at 1.21.

And that is the point, because of Section 14. **The horizon does not charge in elements. It charges in barrel.** A thinned array at the deepest rung saves 4.3e5 in elements and pays 9.6 times the barrel, which is 9.6 times the temperature — 4.7e6 becomes 4.5e7. In a real telescope elements cost money and baseline costs land, so an engineer buys sparse every time. Here the closed form already made elements free, three exponentials whatever the node count, while barrel is the size of the integers one must reach and barrel is exactly what the horizon taxes. In this instrument the **filled** aperture is the budget option, and thinning it is the extravagance: it pays in the scarce currency to save in the abundant one.

14.2 The illumination, and what the constant belongs to

Section 14.1 changed the aperture's geometry. This one changes its illumination, which turns out to matter more.

Redoing the scaling with an arbitrary taper across the barrel gives a profile written entirely in terms of the window's own transform, so any illumination can be swept. Running the uniform window back through that general form returns 1.2087104 against the directly solved value — agreement 3.05e-08 — so the general derivation is checked before it is used.

Swept, a mild taper raises the splitting point to **1.5026** and an edge-weighted illumination lowers it to **1.0919**: the classical trade of resolution against sidelobes, arriving here from arithmetic rather than optics. So the constant is a property of the aperture *and* its illumination, and 1.2087104 is specifically the uniform-illumination value.

The strong windows do something sharper, and it is the most instrument-specific result in this paper. Hann, Hamming and Blackman **never split at any barrel length**. At one barrel, four, eight and twenty the profile carries exactly one maximum, creeping toward the far tooth at 0.415, 0.949, 0.988, 0.998. The reason is that our two teeth are not symmetric: one is a comb's main lobe, while the other is born from interference between the two terms and rides on the window's transform evaluated at the aperture scale. A uniform window leaves that transform decaying slowly; a window designed to suppress sidelobes crushes it, and takes the second tooth with it. Strong apodization here does not blur the pair — **it deletes one of them**. Which is the comma's lesson again: the second tooth is a sidelobe-regime creature, and instruments built to suppress sidelobes suppress it. Measured in the scaled limit; a finite-aperture cross-check of the deletion has not been run. [measured / open]

One gap we report as a gap. Our two teeth are equally bright throughout, and a weak line hiding under a strong line's sidelobe is the first thing a spectroscopist asks about. We built that test and **rejected it**: every amplitude ratio down to one twentieth returned the identical barrel, which is not a result but a broken test. Near the kiss both combs have a tooth nearby, so a maximum sitting near the weak line can be manufactured entirely by the strong comb's own structure; the test counted the bump and never noticed the dim line was never seen. So we ran the control the test was missing: delete the weak comb entirely and look again. At 1.25 barrels a peak still sits beside the weak tooth **with the weak comb gone** — 108.6496 against 108.6525 — which settles it: the original test was reading the strong comb's own sidelobe. With attribution enforced, requiring the peak to be present with the weak comb and absent without it, the weak line is genuinely seen at four and eight barrels and not before.

What we still cannot give is a clean price. The corrected test returns barrels in steps rather than a curve — 1.75 down to a ratio of 0.3, then 3.00 down to 0.02, then nothing at 0.01 — and the steps track the prominence floor we chose, which is the same disease as the tolerance that produced our wrong constant. So: the attribution question is answered, the pricing question is not, and the constant reported here remains an equal-amplitude value. [attribution measured; price open]

15. What is ours, and what is not

The criterion is Sparrow's, from 1916. That coherent sources behave differently from incoherent ones under it has been known since Barakat's work in the early 1960s, and the modern line — two points with unequal amplitudes and a nontrivial relative phase — is active today in spectral super-resolution. The floor beneath every coincidence here is Baker's. The bargain distribution is Gauss and Kuzmin's. The primon gas is Julia's. The comma's identification with a Baker-floor row is Barbour's, from 1948, as the Crossover Atlas established when it downgraded its own verdict.

Barakat's apodization papers solve the *inverse* problem — which pupil function optimises the Sparrow limit, by calculus of variations and a Fredholm equation — rather than tabulating limits for standard windows, so the sweep in Section 14.2 does not duplicate them.

For comparison we computed the classical slit-aperture Sparrow limits in the same units: incoherent gives 0.8296, two coherent sources in quadrature also 0.8296, in-phase 1.3252, and counter-phase splits at essentially any separation. Ours is none of them, because our profile is not two point-spread functions added together — it is a shifted comb minus a constant, divided by the distance from one tooth, a shape that exists because the aperture is a lattice built from two prime towers rather than a slit.

We could not locate our constant in the literature. A numerical hunt for a constant is weak evidence, and we swept the optics and super-resolution side without touching the number-theoretic side, where someone may have reached this wall from the other direction.

What we claim is the instrument and the census, which are ours by construction, and two results that we believe are ours until somebody shows otherwise.

The first is the constant. It is not located in the literature, and "not located" is the strongest thing an uncredentialed search entitles anyone to say — but it is also not nothing. The number is solved rather than fitted, to twenty digits, and it is confirmed by 13,978 independent measurements across 993 arithmetic worlds and by giant-prime rungs with ten-billion-column apertures. If it exists elsewhere, it exists under a normalisation nobody has pointed us to.

The second is the horizon, and here we must correct an assessment we made earlier in this work and got backwards. It is tempting to rate a wall at exactly π as the more exotic of the two results. It is the less exotic. Its π is the radian convention inherited from the aperture identity, its nontrivial content is the number one half, and the algebra that produces it — two side maxima of a reciprocal product of two quadratics survive precisely while the half-width stays under half the separation — is elementary enough that somebody has certainly written it down. We would be unsurprised to be shown it.

What is ours in the horizon is not the number but the *identification*: that the half-width in question is the gas temperature multiplied by the barrel the coincidence demands, divided by a turn. That a thermodynamic parameter sets a resolution wall on an arithmetic aperture is the claim, and the wall's location follows from it by school algebra.

By that standard the constant, not the horizon, is the stronger novelty candidate. It is the first root of a transcendental scalar function; it does not simplify against π , e , or any small rational we tried; and it survived a test built to kill it at primes eleven digits long. If either of these two results turns out to be genuinely new, our present expectation is that it is that one.

Nobody in 1916 was aiming Sparrow's criterion at prime towers.

16. How to break this

Find a patch whose deep rungs do not split at the solved value, and universality dies. Find the constant already in the literature under another name, and it is not ours — we would far rather cite you than be corrected later. Identify it in closed form, which we could not: it is the first root of an explicit function of one variable and should not be hard for someone who recognizes the shape. Or show that the closed form's neglected term matters at some depth, in which case the deep half of the census needs redoing.

17. What this paper does not do

Three receipts belonging to this program remain proposed and unbuilt, and we name them rather than let them pass: a zero count against a winding count; the functional equation measured as a fold; and a ribbon's twist and writhe trading while the link number holds integer.

Nothing here claims new physics. The instrument is ours; every theorem belongs to its people.

Methods and receipts

Every figure and number regenerates from code shipped alongside this paper. Each tool writes a stamped ledger containing every number it produced, and no number appears here that does not appear in one.

Receipt	What it establishes	Section
magnetic_forest.py	flatness under characters; gauge invariance; flux dependence; butterfly	1, 2
rotor_screen.py v1.1	open-beam flatness; combs and the comma; the silences; flux screen	3
comma_vernier.py	the comma as rulers, dial and ladder	4
thermal_forest.py	the Bose backbone; the Hagedorn knee; apodization; the horizon; the current	5-8, 14
comma_resolver.py v1.1	the split measured against the honest node sum	9, 10
closed_form_beam.py v1.1	the sum-free beam, the collapse, the solved constant, the deepest rung	10, 11
patch_census.py v1.1	993 patches, ladder recovery, the bargain census, 13,978 measured rungs	12, 13
clone_ledger.py	the clone rhythm, the gauge phase, the lost barrel, census completeness	12.3
giant_pair_test.py	the falsification test at giant primes, with its precision boundary declared	12.4
sparse_array.py	the redundancy census and the two-currency trade	14.1
aperture_variations.py	the illumination sweep, and the rejected unequal-amplitude test	14.2

Runs are deterministic; no random number generator affects any reported value. The census run reported here surveyed 993 of 1,035 available patches in 48.8 minutes; the remaining 42 are the largest-prime pairs and are required by no claim.

Tag register

[exact] — algebra, verified to machine precision: the flatness results, the aperture identity, the Bose backbone, and the π horizon. **[measured]** — produced by a shipped receipt: the flux response, the comma gap, the splitting thresholds, the Hagedorn knee, the census, the constant, the persistent current. **[kinship, fenced]** — structural resemblance without a claim of identity: the butterfly and Hofstadter, the beam's dips near the low zeta zeros, the current and mesoscopic rings. **[open]** — stated and not established: the closed form of the constant, the combination of the horizon with the full gas's temperature cap, the horizon's status as a peak-counting rather than information-theoretic limit, the price of unequal line strengths, the finite-aperture behaviour of the tooth deletion, and everything in Section 17.

References

- Sparrow, C. M. (1916). On spectroscopic resolving power. *Astrophysical Journal*, 44, 76.
- Barakat, R. (1962). Application of apodization to increase two-point resolution by the Sparrow criterion. I. Coherent illumination. *Journal of the Optical Society of America*, 52(3), 276-283.
- Barakat, R., & Levin, E. (1963). Application of apodization to increase two-point resolution by the Sparrow criterion. II. Incoherent illumination. *Journal of the Optical Society of America*, 53(2), 274-282.
- Baker, A. (1966-1968). Linear forms in the logarithms of algebraic numbers, I-III. *Mathematika*.
- Barbour, J. M. (1948). Music and ternary continued fractions. *American Mathematical Monthly*, 55(9), 545-555.
- Hofstadter, D. R. (1976). Energy levels and wave functions of Bloch electrons in rational and irrational magnetic fields. *Physical Review B*, 14, 2239.
- Julia, B. (1990). Statistical theory of numbers. In *Number Theory and Physics*, Springer.
- Khinchin, A. Ya. (1964). *Continued Fractions*. University of Chicago Press.

18. Tool and computational resource disclosure

This disclosure belongs to The Magnetic Forest: a two-channel instrument for the integer orchestra, the gas it turns out to be, and the price of a coincidence, concept DOI 10.5281/zenodo.22077658, ORCID 0009-0000-0956-5743. It is the final page of that record.

Tools used. Claude (Anthropic) was used for prose drafting, for authoring the eleven receipt scripts deposited with this record, for literature search, and for the build and checking tools described below. Every numerical result reported in this paper was produced by those deposited scripts, each of which writes a timestamped ledger containing every value it computed; no reported number was transcribed from a conversation. Typesetting used python-markdown and wkhtmltopdf. Symbolic verification of one algebraic factorisation used SymPy, and one constant was solved at forty digits with mpmath.

Responsibility. Responsibility for the correctness and adequacy of the arguments and results presented here, and for the completeness and accuracy of the citations, rests entirely with the named human author. No automated system is an author of this work.

Verification performed. Every reported number is regenerated by the deposited code rather than transcribed, and each figure carries the tool name and run stamp of the ledger that produced it. The splitting constant was solved rather than sampled after a sampled estimate was found to drift in its seventh figure, and the solved value was independently recovered by a second route in the thermal receipt. The closed-form beam was validated against the exact node sum at every depth where that sum can still be evaluated, and the boundary beyond which no independent cross-check exists is declared in the paper rather than passed over. Two tests were built, run, found to be measuring their own artifacts, and are reported alongside their corrections rather than deleted. Citations were checked against the publisher record. The final artifact was checked with lint_deposit.py, with a build-time font preflight asserting that every non-ASCII codepoint is covered by a named embedded font, a margin assertion measured from the built page rather than requested of the renderer, and a check that no heading falls as the last line of a page.

Attribution limits. Two prior-art searches were run: one across the optics and spectral super-resolution literature, and one across the apodization literature. Neither located the splitting constant reported here, and novelty language in this paper is accordingly restricted to "not located" and never to "new." Three specific gaps are stated rather than left implicit. No credentialed database pass (zbMATH or MathSciNet) has been run. The number-theoretic side, where exponential sums over smooth numbers would be the natural place for this result to have been reached from the other direction, was not searched. And a published commentary is known to exist arguing that the standard statement of the Sparrow limit is inadequate for complex pupil functions — the case this paper is in — but its volume and page could not be verified against the publisher record, so it is not cited here rather than cited on an unchecked locator.

Computational resources. All eleven receipts run on a single desktop machine. Ten complete in under a minute each; the prime-pair census reported in Section 12 ran for 48.8 minutes over 993 prime pairs. No cluster, grid or accelerator was used.