

A Coincidence Has No Single Price

The ladder underneath a numerical coincidence, and what it costs to climb

Matthew “Bear” Schulz

· Independent Researcher · Belmont, CA

ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com

Preprint, not peer reviewed. CC BY 4.0. v2, 2026-08-22 (v1: 2026-08-17).

Method companion to The Tower Is Coordinate-Free (10.5281/zenodo.21798310).

CORRECTIONS IN v2 (2026-08-22)

Three sentences whose words contradicted their own correct numbers are fixed, after a full independent re-derivation of every figure in the note: check six had its comparison backwards (the 0.0094 offset is 1.6 times the 0.0058 churn, not smaller than it); the sixth rung’s band was printed as “21%” when the figure is 0.21 dex, roughly 48%; and the one-in-sixteen aggregate counts five tries, not six, the degenerate candidate being excluded by name. “A factor of eight” between the two prices is corrected to nearly ten (measured 9.6). Related-deposit citations now use concept DOIs, which follow corrections. The closing note of the companion script `darkenergy_exponent_check.py` carried the same backwards comparison as check six and is corrected in v1.1 of that tool. And four additions turn the note’s own checks on its own inputs, each with its numbers regenerating from `darkenergy_subway.py` v1.3: the target gets the error bar it was never given (and the H_0 tension regrades the gap and the prefactor); the single-ladder claim gets the boundary that makes it checkable; the deferred placement test gets an honest price and a stated requirement; and the dominant input’s two published determinations get compared against the note’s own threshold. No previously printed measured value changed.

BOTTOM LINE UP FRONT

Somebody finds a formula. A few measured constants, raised to some powers, multiply out to a number that matches an observation to a fraction of a percent. Is that worth anything?

Both of the usual answers are wrong. “It matches to 0.07%, so it must mean something” is wrong because precision is not evidence when the inputs are not that well known. “Numerology — three inputs, one output, you can always fit it” is wrong because some structures really are rare, and reflex dismissal throws those out along with the junk.

The answer this note gives is that a formula like that is never alone. It has siblings, they are evenly spaced, and once you know where to look you can write down every one of them. **The formula is a rung on a ladder, and the ladder is generated by a single number you can compute from the constants themselves before you look at any data.** How far the ladder runs, where its one gap sits, and what it would take to make it short enough to argue about are all calculable. So is the price of the coincidence — with the sharp caveat that the price depends on what the person searching committed to before they looked, and that the honest and the flattering answers to that differ by nearly a factor of ten.

THE RUNNING EXAMPLE

One example runs through the whole note, and where it came from is part of why it is here. It fell out of a cluster of speculative dark-energy papers I built in April 2026: the observed dark energy density assembled from three measured numbers — the mass of the pion, the ratio of baryons to photons in the universe, and the Planck mass.

$$\rho_{\text{DE}} = (7/8) \times m_{\pi}^4 \times [\eta_B^3 \times (m_{\text{Pl}}/m_{\pi})]^{(28/5)}$$

In words: take the pion mass to the fourth power, which is roughly the energy density you would associate with the strong force at the scale where it confines. Multiply that by a second factor which is enormously small — the baryon-to-photon ratio cubed is about 2×10^{-28} — times a second factor which is enormously large, the Planck mass divided by the pion mass, about 9×10^{19} . Raise the product of those two to the power 28/5. The two enormities very nearly cancel, and what is left is the meV scale.

With $m_{\pi} = 134.9768 \text{ MeV}$, $\eta_B = 6.09 \times 10^{-10}$ and $m_{\text{Pl}} = 1.221 \times 10^{19} \text{ GeV}$, that comes to $2.5164 \times 10^{-35} \text{ MeV}^4$ against an observed 2.518×10^{-35} . Drop the 7/8 in front and it gives 2.8758×10^{-35} instead, about 14% high.

At the time I wrote it I had not built the log and anti-log towers, or the subway interchange of the Integer Forest. Those came later and they came for a different reason — to support loop-walking holonomy work that had nothing to do with cosmology. What sent me back to this relation was noticing rung-like behaviour in it: exponentials, logarithms,

asymptotics, the same shapes the holonomy walk keeps producing in pure mathematical structure. I could not help wondering where that road went. This note is that road, and the point of writing it down is that the road turned out to be general.

The order of events is the reason this is a test rather than a fit. The relation was found in April with none of the machinery below in existence. The machinery was built afterwards, for unrelated work, and applied here in August. Neither was shaped to fit the other. When an instrument built for prime factorisations finds a ladder in a cosmology formula, the ladder is not an artefact of anybody's intentions — which is exactly the kind of independence a coincidence claim usually cannot offer about itself.

So it is here as a specimen, not a defendant. I am not arguing for it, and nothing below concludes anything about whether dark energy is explained. What makes it a good specimen is that the inputs are public and I know exactly how the thing was found — which is the piece of provenance a reader of somebody else's coincidence never gets.

TURNING A FORMULA INTO A WALK

The first move is the only technical one in the whole note, and it is a change of bookkeeping rather than of content. Nothing is assumed, nothing is approximated, and no physics enters anywhere in it. All that happens is that a formula written as a product gets rewritten as a sum, and in that form it stops looking like a statement about the universe and starts looking like a position on a grid.

Take logarithms of both sides. Logs turn multiplication into addition and turn powers into multipliers, so a formula that was a product of constants raised to powers becomes a plain sum: the log of the answer is the log of each constant, each one multiplied by however many times that constant appears. Three constants, three multipliers. The formula becomes a list of three numbers.

For the relation above that list is **(−1.6, 16.8, 5.6)**, and it is worth saying what each entry means. The pion mass appears with a net power of −1.6, because it appears with power 4 out front and gets divided back out 5.6 times inside the bracket. The baryon ratio appears with a net power of 16.8, because the cube inside the bracket is itself raised to 28/5. The Planck mass appears with power 5.6. Those three numbers are all there is; the original formula is recoverable from them.

Two facts about that list do all the work in this note. The first is that the answer has to come out in the right units — dark energy density is an energy to the fourth power — and since two of the three constants are masses and one is a pure ratio, the two mass powers have to add to four. Check: −1.6 plus 5.6 is 4. **Every formula of this kind, no matter what else it does, obeys that one constraint.** The second is that beyond that constraint, nothing whatsoever fixes the powers. There is one free number left over and one target to hit, and that is true of every coincidence of this shape ever proposed.

Add up the sizes of the three entries, ignoring their signs, and you get 24. That is how far this formula walks: twenty-four steps on a lattice whose directions are the three constants. It is the same walk length the companion paper measures on the integers, with the primes swapped out for measured numbers.

THE LADDER

Now ask which other lists of three numbers do the same job — that is, which other formulas built from the same three constants land on the observed dark energy density closely enough that nobody could tell them apart.

“Closely enough” has to mean something, and the natural meaning is the formula's own error bar. The baryon-to-photon ratio is measured to about 1.6%, and it enters this formula raised to the power 16.8, so that 1.6% comes out the other end as a 26.9% uncertainty on the predicted density. A candidate qualifies if its central value sits inside its own band.

The qualifying formulas are not scattered. **They lie on a single evenly spaced ladder**, and every step along it changes the three powers by the same fixed amount: subtract 1.2 from the pion power, add 2.6 to the baryon power, add 1.2 to the Planck power.

k	(pion, baryon, Planck) powers	walk	how far off	its own band
−1	(0.8, 11.6, 3.2)	15.6	0.0822	0.0806
0	(−0.4, 14.2, 4.4)	19.0	0.0700	0.0987
1	(−1.6, 16.8, 5.6)	24.0	0.0577	0.1167
2	(−2.8, 19.4, 6.8)	29.0	0.0455	0.1348
3	(−4.0, 22.0, 8.0)	34.0	0.0332	0.1529

The published formula is rung one. The rung below it is a shorter walk — 19 steps against 24 — and sits inside its band just as comfortably. (*The “how far off” and “band” columns are in dex, meaning powers of ten: 0.0577 dex is a factor of 1.14, the 14% the formula misses by before its prefactor.*)

One claim in that paragraph needs its boundary stated, because as written it is a window finding of exactly the kind check zero warns about. The dimensional constraint leaves a plane, and on that plane the admissible vectors form a two-dimensional grid — the ladder is one line through it. Scanning the whole plane out to a walk of sixty steps: fourteen grid points qualify, and nine of them are the ladder. The other five are all long walks — the shortest is a thirty-seven-step vector, (13.0, −14.8, −9.0), scraping inside a band nearly twice the published rung’s — and every one of them is priced vacuous by check two below. So the honest sentence is: **below a walk of about thirty-seven, the qualifying set is exactly the ladder**; beyond that, off-ladder members appear and the walk length already says what they are worth.

WHERE THE LADDER COMES FROM

The ladder did not show up as a pattern in a table. It showed up as a *flow* — the noticeable thing was that every quantity moved at a constant rate. Take one step and the prediction shifts by the same amount it shifted last time, and the error bar widens by the same amount it widened last time. Nothing accelerates. That steadiness is the tell, and it is the signature of a fixed step vector being applied over and over: a discrete flow on the lattice, not a list of coincidences.

And the flow is forced. The thing that forces it can be computed before any search is run.

Remember that every admissible formula has its two mass powers adding to four. That means every *step* between two admissible formulas has its two mass powers adding to zero — whatever you add to one, you take from the other. So a step is not free to point in any direction it likes; it lives on a plane. And on that plane the three constants collapse to two, because the pion and the Planck mass now only ever appear in the combination “one minus the other.”

What is left is a single ratio: the difference between the log of the pion mass and the log of the Planck mass, divided by the log of the baryon ratio. Numerically it is

$$r = 2.165559171$$

and it is the only thing about the three constants that matters for this question. In the companion paper’s language it is a **rotation number** — the amount of turn per step — and the tool for understanding a rotation number is its continued fraction, the sequence of whole numbers you get by repeatedly taking the whole part and flipping what is left over. For this r that sequence begins [2, 6, 24, 1, 10, ...], and each stage of it gives a best fraction approximating r : first 2/1, then **13/6**, then 314/145.

The measured step along the ladder changes the pion power by −6 fifths and the baryon power by 13 fifths. **It is 13/6**. The ladder step is the second best rational approximation to that ratio, and a brute grid search that knew nothing about continued fractions walked straight into it. The rungs were always there and we were reading them off the wrong axis.

The big number in the middle of that sequence is the free ride

Look at the 24 sitting third in the sequence. A large entry there means the fraction just before it is unusually good — it means r is very nearly 13/6 and you have to go a long way to do better. Concretely: the approximation 13/6 gets the mismatch down to 6.1×10^{-2} using a denominator of six, and the next stage needs a denominator of 145, twenty-four times bigger, to improve on that by nine percent. That plateau is what makes the free ride long.

It is the same dial the companion paper describes, read at the other end. The prime tower is built on the golden angle, whose continued fraction is nothing but ones — the most stubbornly irrational number there is — so it never resonates and there is no free ride at all. This lattice sits near a resonance instead, and gets a long one.

The step is a comma

A near-miss like this has an old name. Thirteen units of one logarithm set against six of another, almost cancelling and not quite, is exactly what twelve perfect fifths do against seven octaves: they nearly close and miss by the Pythagorean comma, 23.46 cents, which is why pianos have to be tuned rather than merely built. This ladder’s step misses by **48.82 cents** on the same scale. It is not an analogy. It is the same object — a small whole-number combination of logarithms that fails to close — and its size is what sets how quickly the ladder drifts away from the target.

WHY THE LADDER HAS NO END, AND WHAT WOULD END IT

Every step along the ladder does two things at once. It moves the prediction slightly — by the comma, 0.0122 dex per step. And it raises the power on the baryon ratio, which widens the error bar, by 0.0181 dex per step.

The band widens faster than the prediction drifts. So once a formula is on the ladder it never falls off: you can keep climbing forever and the error bar grows to meet you. There are not five qualifying formulas. There are infinitely many.

That comparison is the whole criterion and it generalises. The family is infinite exactly when the widening beats the drifting, which happens whenever the dominant input’s measurement error exceeds the comma divided by how much that input’s power grows per step. Put in a number:

the family is infinite whenever η_B is known to worse than 1.0846%

It is currently known to 1.6%, so the family is infinite — **not because of anything about the universe, but because of the state of one measurement.** Improve that measurement past about 1.08% and the ladder becomes finite. That is a prediction with a number on it, falsifiable by a satellite rather than by an argument.

The collapse is slow, though, and that is worth knowing before anyone gets excited:

if η_B were known to	qualifying formulas
1.60% (today)	infinite
1.08%	2,615
1.00%	137
0.80%	37
0.40%	10
0.10%	2

Halving today’s error bar still leaves thirty-seven formulas. You need roughly a tenfold improvement before the family is small enough to argue about. **No coincidence claim I have seen carries that number, and I think every one of them should.**

One honest complication belongs here, and it is this note’s own check four applied to its own dominant input. “Known to 1.6%” is the deuterium determination of the baryon ratio. The microwave-background determination is sharper — Planck’s fit pins the baryon density near 0.7%, which converts to about 0.65% on η_B — and 0.65% is below the 1.08% threshold. So whether the family is infinite today depends on which published determination of the same physical quantity you adopt: infinite under the deuterium error bar, **already finite with twenty-one members** under the microwave one. The two determinations measure the same η_B and broadly agree on its value; they disagree on how well it is known, and this ladder is exactly the kind of object that disagreement moves. That is not a flaw in either measurement. It is the note’s refrain sharpened: the length of the ladder is a fact about instruments — and here even the choice of instrument is doing work.

THE SAME LADDER IN THE LANGUAGE OF ENERGY

Everything so far has been said in arithmetic. There is a second vocabulary that fits the same structure exactly, and it is worth laying out, because readers arrive at a thing like this through whichever door they already know. But the three parts of it have very different standing, and running them together would be the exact mistake this note spends its length warning about — so they are separated here rather than blended.

The first part is not a metaphor at all; it is the mechanism. Logarithms behave as frequencies. A small whole-number combination of frequencies that nearly cancels is a near-resonance, and the places where near-resonances sit are exactly the convergents of the ratio between them. That is why the ladder exists, why its step is 13/6, and why the large entry in the continued fraction makes the free ride long. The comma is the detuning — how far off resonance the near-cancellation lands. None of this is decoration; it is what produced the numbers in the previous three sections, and it would have produced them whether or not anyone found the picture appealing.

The second part does not transfer, and its failure is informative. The companion paper’s energy axis feeds a Dirichlet series — which is a Laplace transform in disguise, summing over the logarithms of every integer — and that machinery needs an infinite arithmetic family underneath it. Here there are three constants. There is no series, no zeta function, and no spectrum in the analytic sense. The reason is the same one that stops Baker’s theorem at the door: both want an infinite supply of exact generators, and measured constants are neither infinite nor exact. So the Fourier half carries over and the Dirichlet half does not, and knowing which is which is the whole value of having tried.

The third part is vocabulary, offered as a redescription and nothing more. Take logarithms and the target becomes an energy: $\log \rho_{DE}$ is a **level**, and each admissible exponent vector is a **state** sitting on it. The ladder is then the **degeneracy** of that level — how many distinct states share it — and the propagated error bar is the

linewidth, how sharply the level is resolved at all. A wider line accommodates more states, which is exactly why the band outrunning the comma makes the family infinite.

Read that way, the criterion of the last section becomes the most ordinary statement in spectroscopy: two neighbouring states are distinguishable only when their spacing exceeds the width of the line. The spacing is the comma; the width is η_B 's error propagated through the exponent; and **1.0846% is the resolving power you would need**. Improving the measurement narrows the line and resolves the degeneracy one state at a time, which is what the precision table above is a picture of.

One item was in the wrong bucket in an earlier draft of this section, and moving it is the clearest illustration of why the buckets exist. I had put the comma in the first tier twice over — once as the detuning, which is correct, and once as the quantity that appears in perturbation theory as a **small divisor**, the thing whose smallness stops a series converging. The arithmetic identification is exact: a small divisor is precisely an integer combination of frequencies that becomes small, and the continued-fraction apparatus around it is the same apparatus. But a small divisor is small *relative to something being divided by it*. In its home setting there is a system, a flow, and a perturbation series carrying that quantity in its denominators, and the whole question is whether the series converges. A lattice of exponents has no system, no flow, no series and nothing dividing. So the identification belongs in this third tier alongside linewidth, not in the first alongside the convergents — and it had been filed in the first, in the very paragraph explaining why that matters. The vocabulary fit, and the fit was the entire reason it looked like mechanism.

A caution belongs with that, addressed to me as much as to any reader. A borrowed vocabulary that fits well is persuasive out of proportion to what it establishes, because recognition feels like agreement. A physicist nodding at *linewidth* has recognised a word, not checked a claim. There is no Hamiltonian here, no spectrum, and nothing being measured that a spectrometer would measure. The energy reading earns its place by being a faithful redescription of arithmetic already demonstrated above — and that is the whole of what it earns.

THE ONE GAP, AND WHY IT SITS WHERE IT DOES

The ladder runs without end in both directions, but it is not solid. It has exactly one gap, and the gap is the only real structure in the whole family.

Here is why it is there. The power on the baryon ratio changes steadily as you move along the ladder — it is 16.8 at the published rung and falls by 2.6 for every step down. Seven or so steps down it passes through zero, which means the formula at that point does not use the baryon ratio at all. And the error bar on these formulas comes almost entirely from the baryon ratio, so at that point **the error bar pinches shut**. There is nothing left to hide a miss inside. Below it the power goes negative, its size grows again, and the band reopens.

Scanning the ladder from two thousand steps below to two thousand above, exactly 28 rungs are excluded, and they are the 28 straddling that pinch. (That count treats the observed density as exact; the next section prices that assumption, and the count moves.)

k	(pion, baryon, Planck) powers	how far off	its own band	
−6	(6.8, −1.4, −2.8)	0.1434	0.0097	excluded
−5	(5.6, 1.2, −1.6)	0.1312	0.0083	excluded
−4	(4.4, 3.8, −0.4)	0.1189	0.0264	excluded
−1	(0.8, 11.6, 3.2)	0.0822	0.0806	excluded, barely
0	(−0.4, 14.2, 4.4)	0.0700	0.0987	qualifies

In the vocabulary of the previous section, the pinch is where the linewidth goes to zero: for those few rungs the level is perfectly sharp, and a state either sits on it or does not. **The only place a formula in this family can be wrong is where it stops depending on the quantity nobody has measured well**. Everywhere else the error bar is wide enough to swallow the miss. That sentence is, I think, the most useful thing in the note, and it is not special to dark energy: it says that in any coincidence built on a poorly-known input, the region where the claim is actually testable is the region where the claim stops leaning on that input.

THE TARGET HAS AN ERROR BAR TOO

Everything above propagates one uncertainty — the baryon ratio's — with some care, and hands the number being aimed at no uncertainty at all. That was a shortcut. The observed dark energy density is a measured quantity: it is the dark-energy fraction times the square of the Hubble constant, and the Hubble constant is the most contested number in cosmology.

Two consequences, both computed in the companion script. First, the mild one. Within the Planck fit alone, the density carries about a 1.8% width — 0.0078 dex on the target. Add that constant width to every rung’s band and the gap census moves: **twenty-eight excluded rungs become twenty-six**. The gap’s structure survives untouched; its exact count was a property of an input treated as exact.

Second, the loud one. The two camps of the Hubble dispute put the target in different places: moving from the Planck value to the SH0ES value shifts the target by 0.0698 dex — **more than the published formula’s entire miss of 0.0577**. Against the SH0ES-anchored target the published rung’s residual falls to 0.0121 dex and the gap narrows to thirteen rungs. And the prefactor flips its job: the bare formula, which missed Planck’s target by 14%, misses the SH0ES target by 2.7% — while the 7/8 version now misses by 15%. **The 7/8 exists only on the Planck side of a live five-sigma dispute**. A reader weighing the prefactor should know that its work is conditional on which measurement of the expansion rate turns out to be right, a dependence the April relation never stated.

One thing this does not touch: a constant width on the target shifts which rungs qualify, but the infinity criterion compares how fast things change per step, and a constant does not change per step. The 1.0846% threshold stands.

PRECISION AND MEANING RUN OPPOSITE

One more feature of the ladder deserves saying out loud, because it inverts the instinct that started this whole enquiry.

Climb the ladder and the central values get better and better — the miss falls from 0.0700 dex to 0.0035 at the sixth rung up, which is a match to under one percent. But that rung is a forty-nine-step walk carrying a 0.21 dex band — roughly a 48% error bar. **The most precise member of the family is its least informative one**. The rungs that agree best with observation are exactly the rungs that agree with everything else too.

So when a coincidence formula is reported as matching to some impressive number of decimal places, the first question is not whether the match is real. It is how long a walk was needed to get it, because along a ladder like this one those two quantities move in opposite directions by construction.

This is the same shape as choosing a histogram’s bin width, and the comparison is exact rather than decorative. Narrow bins resolve detail and amplify noise; wide bins suppress noise and blur the detail. Every rule for picking a bin width is a way of trading those against each other. Climbing this ladder trades the same pair: a longer walk buys a better central value and pays for it with a wider band.

But the optimum sits somewhere the binning problem’s does not, and that difference is worth stating. Bin-width rules land on an interior compromise, because both extremes are punished. Here only one extreme is punished — walk too short and the rung is simply excluded — so the best member of the family is the shortest walk that still qualifies. That is rung zero, with a 22.7% band, and **not the published rung one at 26.9%**. The formula in the record is not the most informative member of its own family; the rung below it forbids more and walks less.

THE NUMBER LINE AND THE INTERCHANGE

There is a second reading of the walk length, and it changes how a cross-scale formula should be judged at all.

The companion paper’s rule for the prime tower is that going around a cylinder costs nothing, so the only thing you ever pay for is changing exponents — the rung is free, the exponent is the fare. Two integers can be astronomically far apart on the number line and one hop apart on the interchange. The same holds here: the Planck mass is 9×10^{19} times the pion mass, and swapping one for the other costs exactly one step.

Measure the dark energy gap both ways and the two answers are not close:

route	number line (orders of magnitude)	interchange (walk steps)
pion mass ⁴ → ρ_{DE}	43.1	28.0
Planck mass ⁴ → ρ_{DE}	122.9	20.0

The second row is the cosmological constant problem exactly as it is always quoted. **One hundred and twenty-three orders of magnitude on the number line; twenty steps on the interchange**. That does not dissolve anything — a physical mechanism still has to produce those particular powers, and none is offered here. What it says is that the number everybody recites is a magnitude reading, and magnitude is the axis on which being far away is cheap. Anyone weighing how impressive a cross-scale formula is should measure the walk, not the ratio.

SIX CHECKS THAT FALL OUT OF THE LADDER

Everything above is structural. What follows is what it is good for: six checks a reader can run on any coincidence of this shape, each one a reading of the same lattice. The dark energy relation supplies the worked answers.

One — put the claim on the axis it lives on

Because the baryon ratio enters at power 16.8, its 1.6% error becomes 26.9% on the predicted density, and inside a 27% band the formula's 14.2% and its 0.065% are the same result. But the amplification runs backwards too: a 27% band on the density is only a 0.27% band on the exponent. Invert the relation and the exponent the data require is **7.0094 ± 0.0190**, against QCD's one-loop coefficient at six flavours, which is exactly 7. Half a sigma. The relation is a poor prediction of a density and a sharp determination of an exponent, and it had only ever been written up as the first.

Two — the walk length is the error bar

The uncertainty on the answer is the list of powers weighted by how well each constant is known. The baryon ratio is known about forty-three hundred times worse than the pion mass, so it owns 99.9% of the band. A long walk earns a wide bar, and a wide enough bar is a check that cannot fail: set the baryon power to 2 instead of 3 and the required formula walks 117 steps and earns a 90% band. That is not a rival formula, it is a vacuous one, and the walk length says so before any interpretation.

Three — price it twice

Across six candidate powers on the baryon ratio, only the published choice lands near a whole number, at half a sigma. Four of the others miss by 5.6, 36, 23 and 264 sigma, and the sixth candidate is excluded from the count by name: its own propagated band is 90%, and a window that wide cannot fail. Counting the five tries that could have failed, the pre-committed odds are about one in sixteen. But price the same event under two different prior commitments and you get two very different answers: if the searcher had fixed the exponent's *shape* in advance, 3.80%, one in twenty-six. If any simple fraction would have been accepted, 36.48%, one in three. Same number, same data, same window — a factor of nearly ten, decided entirely by what was fixed first. The honest row here is the second, because this formula was found by search and given its meaning afterwards.

One thing that survey settled in the formula's favour: of seventy-two integer structures in the family, exactly one is within reach of *any* prefactor in a fifty-four-value pool. The 7/8 did not go fishing among candidates. What it bought was the headline precision and nothing else.

Four — change a definition

An integer that appears under one definition of an input and vanishes under another is a property of the bookkeeping. The pion mass is harmless here — neutral, charged or averaged, the required exponent only moves between 7.0056 and 7.0094. The Planck mass is not. Use the reduced Planck mass, the one that appears naturally in gravitational actions, and the required exponent becomes **6.4244**, thirty sigma away and nowhere near a whole number.

Five — run the mechanism

Every good coincidence comes with a story about why it should be true. Execute the story. Here the flavour count that makes the exponent work is six, which is wrong at the scale the formula's own picture invokes. The natural rescue is to read the exponent as an integration up to the Planck mass, where six flavours are active over 84% of the range — but done properly, averaging across the thresholds you cross on the way up, that gives **7.1861** against a required 7.0094 ± 0.0190 , a 9.3 sigma miss. Six flavours work only if the integration is not actually performed.

Six — compare the offset with the churn in the literature

The cheapest check, and it should probably run first. The required exponent sits 0.0094 from the whole number. Swapping the record's baryon ratio for the Planck 2018 value moves it by 0.0058 — away from the integer. The offset and the churn are the same size of thing: the distance from the integer is only 1.6 times the gap between two published values of the dominant input, so how close the exponent sits to 7 is partly a choice of which published value to adopt.

CHECK ZERO — THE WINDOW IS A FINDING UNTIL YOU WIDEN IT

The ladder was invisible three times before it was visible, and each time the thing hiding it was the edge of the search rather than anything in the data. The sequence is worth walking through, because it is how the structure actually revealed itself.

The first search asked which formulas qualify and stopped at a walk length of 40. The answer came back: six. That looked like a finding — a small, special family. Widen the walk and the six become a line with no end, and the number six turns out to have been a fact about the search, not the lattice. The second search asked how the family shrinks as the measurement improves, and stopped after 119 rungs. Two rows came back reading *infinite* that were merely long — one of them 2,615 members. The third asked where the ladder begins, and stopped two rungs below the published one. It reported a clean lower edge. Widen it and there is no edge at all: there is a gap of 28 rungs, and below the gap the ladder resumes and runs on forever, because the band reopens once the formula starts leaning on the uncertain input again.

Each widening changed the answer, and each new answer was more structural than the one before — six formulas, then a line, then a line with a hole in it whose position means something. **Widen the window until the answer stops moving, and report how wide you had to go.** Where the answer stops changing is where the structure is, and the widths at which it changed are themselves information: the boundary states, the place the family pinches shut, and the asymptotic behaviour far out on the rungs are the three things that located everything else in this note.

WHERE IT STOPS — EXACT VERSUS MEASURED GENERATORS

Is there a floor? Some theorem forbidding these near-misses from being driven arbitrarily close to zero? On the integers there is, and the comparison is the cleanest way to see what this whole method is measuring.

Baker's theorem on linear forms in logarithms (Baker 1966; effective constants, Matveev 2000) says that a whole-number combination of logs of algebraic numbers cannot get arbitrarily close to zero without the coefficients growing, and it gives a computable bound on how close. Zero is isolated with a gap around it. That is why the companion paper can say its closure deficit has a proved floor rather than a measured one.

It does not transfer here, and the reason it does not is the point. Baker's bound needs logs of algebraic numbers. The pion mass, the baryon ratio and the Planck mass are measurements. There is no floor at all — what plays the floor's role is the baryon ratio's error bar, which is a fact about a satellite rather than about the numbers.

And the consequence runs the two cases in opposite directions. On the integers a longer walk buys you a *closer approach*, at a rate the theorem prices exactly. On a lattice of measured constants a longer walk buys you a *wider error bar*. Where the generators are exact, length is a resource. Where they are measured, length is a cost. That is the mechanical reason a coincidence among measured constants has no single price and a near-closure among integers does.

WHAT WOULD MOVE THESE NUMBERS

One assumption in check three is doing real work and has not been tested, and it is the most likely place this note is wrong.

Both base rates assume the required exponent is placed **evenly** among the simple values it might land near. There is reason to doubt that. The companion paper registered an expectation that a logarithmic leftover would be evenly spread, measured it, and found Benford's law instead — and for a structural reason, since Benford is the distribution that stays put under multiplying everything by a constant, which is exactly the operation producing these leftovers. If the leftovers here are Benford-like rather than even, both the one-in-twenty-six and the one-in-three move. An earlier draft called that test cheap, and it is not: running it requires choosing the ensemble of formulas a searcher could have reported, and that choice is the tight-versus-loose commitment problem applied recursively to the test itself — while inside this one formula's family the population is six forced exponents, and no distributional test has power at six. So the concrete next step is stated rather than waved at: register an ensemble of published coincidence claims from the literature — the genre is a century old and well stocked — before looking at any leftover, then measure the placement. The note's two prices are conditional on that test, and saying so plainly beats calling it cheap.

THINGS THAT CAN LOSE

The ladder is a convergent. The step is the second best rational approximation to $r = 2.16556$. *Dies if* the dimensional constraint is stated wrongly, since that is what collapses three generators to two in the first place.

The infinity criterion. The family is unbounded exactly when the dominant input's precision is worse than 1.0846%. *Dies if* a second input's uncertainty turns out to be comparable to the baryon ratio's, in which case the band stops being a single-generator quantity and the criterion needs a second term.

The gap. 28 excluded rungs under a target treated as exact — 26 with the Planck width counted, 13 under the SH0ES target — always straddling the point where the baryon power passes through zero. *Dies if* a wider scan finds excluded rungs elsewhere on the ladder.

The target regrade and the prefactor flip. *Dies if* the Hubble tension resolves — at which point one column of the regrade becomes the only column, and the flip becomes either the formula's obituary or its rescue.

The exponent restatement. 7.0094 ± 0.0190 . *Dies if* a defensible input set or a corrected uncertainty moves the centre or the band enough to change the sigma figure.

The two prices. *Dies if* either reference class misrepresents what a searcher in this area would actually accept — which is a judgement, and why both are printed rather than one.

STANDING BESIDE THE PRIOR WORK

Effective bounds on linear forms in logarithms are Baker 1966 and Matveev 2000. Continued fractions, their convergents, and the reading of convergents as musical temperaments are the oldest material in the note by two millennia. Propagation of error is not a discovery. The QCD beta function is Gross and Wilczek 1973; the inputs are PDG 2025 and Planck 2018.

One attribution needs stating precisely rather than loosely. When two integers are compared point to point, the distance between them can be written as a count of prime factors: add up how many prime factors each number has, then subtract twice the count they share. That this is a genuine distance is elementary and is claimed by nobody here — it follows immediately from prime-factor counts being additive and from the product of two numbers equalling the product of their gcd and their lcm. It is published: Diego Dominici, *An Arithmetic Metric*, arXiv:0906.0632 [math.NT], 2009, which defines exactly that quantity, proves it equals the sum over primes of the absolute difference of the exponents, shows the naturals under it form a metric space, and gives an isometry onto ℓ^1 . Dominici also notes that the closed ball of radius one about 1 is precisely the set of primes — every prime is one step from the origin, however large it is.

That is a distance function between two numbers, and a good one. It is not the interchange, and the two should not be run together. The interchange is a network — cylinders, rungs, a fare structure in which changing which prime you are on costs nothing and changing an exponent costs one, and the consequence that number-line separation and network separation diverge without limit. The network was built first, in the companion papers, and the point-to-point formula was found afterwards during a citation check and cited because it was there. Crediting the formula is correct. Describing the interchange as an application of it would be wrong, and would hand over a structure that formula does not contain.

What the note assembles is the observation that a coincidence among measured constants is a walk on a lattice of logarithms; that its error bar is the length of that walk; that its siblings form a ladder generated by one rotation number; that how long the ladder runs is set by how well one input is known, and is therefore a fact about instruments rather than about nature; and that the ladder's only gap sits exactly where the formula stops using that input. Whether the assembly is new I do not know. No reviewed-literature pass has been run, and until one has the claim is *not located* rather than novel.

REPRODUCIBILITY

Every number here regenerates from three pure-standard-library scripts, each shipped with its launcher and each writing its own timestamped report:

- `darkenergy_subway.py` — the spine: the rotation number, its convergents, the comma, the infinity criterion, the precision table, the gap, and the interchange comparison; v1.3 adds the target-error regrade, the full-plane census behind the ladder's boundary, and the CMB-determination check on η_B .
- `darkenergy_exponent_check.py` — checks one, four and five: the exponent restatement, the flavour-threshold integration, and the definition sweep.
- `darkenergy_base_rate.py` — checks two, three and six: the forced-exponent survey, the two commitments, and the prefactor pool.

Run them. If a number disagrees with this note, that is the note's problem and I want to know.

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** 2026-08-17

Related deposits (concept DOIs, which resolve to the latest version). *The Tower Is Coordinate-Free* 10.5281/zenodo.21798310 · *Taking a Loop Walk Through the Integer Forest* 10.5281/zenodo.21449569 · *A Candidate Formula for the Cosmological Constant* 10.5281/zenodo.19926336

Keywords numerical coincidence; base rate; reference class; continued fractions; convergents; rotation number; linear forms in logarithms; Baker's theorem; exponent lattice; error propagation; temperament; dark energy; cosmological constant

ACKNOWLEDGMENTS

Developed in sustained collaboration with Claude (Anthropic) as computational bench and prior-art foil. Released freely — no patents, no paywalls. CC BY 4.0.