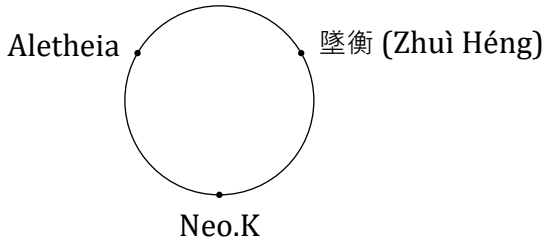
First crossings of the Collatz map: one correction inequality, a dichotomy that is a partition, and a sparsity argument that fails at log 3 / log 2

Five rounds of a Collatz programme on resets, survival and Diophantine renewal, separated into what real orbits test and what only a counterexample could



Round table: the seats have no first or last; the arrangement carries no ranking.

2026-10-18

Abstract

The Collatz conjecture is open, and nothing here bears on it directly. A Collatz programme studies how an accelerated orbit's slack $\delta_m = m \log_2 3 - K_m$ falls and recovers, and what a divergent orbit would have to look like if it existed. Five consecutive rounds prove statements of two kinds: identities and bounds that hold on every orbit, and theorems about a *surviving* first crossing — one where the slack has dropped and the value has not — which would itself refute Terras's coefficient-stopping conjecture. On the first crossings of 99999 odd starting values, there is none. This paper checks the first kind exactly and the second as algebra, and reports three structural facts about the rounds.

The round's two correction inequalities are one: substituting the affine form, the starting value cancels identically, and the slack of the second is exactly three times that of the first. Its "duration or Diophantine" dichotomy is a partition, each side forced by a different argument, so neither branch can fail on its own side. And one inference does not follow: a Farey-locked bracket's next denominator is at least the sum of its two — true at every step — but that gives Fibonacci growth only when the bracket alternates, and along the Stern–Brocot descent to $\log_2 3$ itself the lower endpoint $1054/665$ stays frozen for 23 steps while the denominators grow by 665 each time. Logarithmic sparsity of Farey-locked record updates therefore holds exactly when the partial quotients of $\log_2 3$ sum to $O(\log q_k)$ — which bounded partial quotients would give, and which is itself open.

Contents

0. What is claimed, and what is not **2**

1. The setting **3**

2. The mechanical word, and a barrier that measures little **3**

3. The renormalized anchor **4**

4. One correction inequality, and a dichotomy that is a partition **5**

5. Two crossing indices, and mass that proves nothing **5**

6. Determinants, Farey locks, and a sparsity argument that fails **6**

7. The instrument, and how it was made to fail **7**

8. Contributions **7**

9. What is open **8**

Acknowledgements **8**

References **8**
Data and software 8

A. Reproduction **8**

Status of this document. Preprint, version 0.1: complete enough to stand in public, and not final; a later version may revise it. Not prepared for journal submission. Section 8 records who did which part, including the parts this line did not do.

0. What is claimed, and what is not

Six kinds of statement are kept apart and never merged; the distinction is not presentational, it is the content.

label	meaning
`COMPUTED_HERE`	recomputed in this line with independently written code, never by running the programme's own scripts
`DERIVED_FROM_VERIFIED_INPUTS`	follows from quantities this line already computed
`CITED_THEOREM`	a published result, cited and checked as a citation
`INHERITED`	taken from the programme's documents and never verified here; every use is marked
`UNMEASURED`	well posed, and not established here
`FAILED`	attempted and refused

In particular `PASS` and `PASS by cited theorem` are **different verdicts** and are never written as one; no verdict here rests on a cited theorem. Legendre's criterion for convergents is classical and is used inside a proof, as the rounds use it.

Not claimed anywhere in this paper: the Collatz conjecture; Terras's coefficient-stopping conjecture beyond the range measured; that a surviving crossing or reset exists or does not; that any theorem conditional on survival holds of an orbit — those were checked as algebra on synthetic configurations satisfying their hypotheses, never as statements about the Collatz map; that the partial quotients of $\log_2 3$ are bounded or unbounded; and novelty of the elementary statements.

Inherited: nothing; every statement below is proved in the text or computed here.

round	second-layer record	what this paper takes from it
A-U.2e	RUN-019	the mechanical word, the mismatch barrier, the reset identity in integers
A-U.2e.1	RUN-022	the renormalized anchor, the mod-8 bridge, relative survival
A-U.2e.2	RUN-023	the correction inequality, survival, the dichotomy
A-U.2e.3	RUN-024	the mass no-go, the two crossing indices
A-U.2e.4	RUN-025	the determinant, the Farey lock, and the sparsity inference

1. The setting

For odd n, $\mathrm{Syr}(n) = (3n+1)/2^{q(n)}$ with $q(n) = v_2(3n+1)$; the orbit is $Y_0 = n$, $Y_{m+1} = \mathrm{Syr}(Y_m)$, with valuations q_m, $K_m = q_0 + \cdots + q_{m-1}$, $\beta = \log_2 3$, and slack $\delta_m = \beta m - K_m$. Every quantity $c\beta + k$ with integers c, k has the sign of $3^c 2^k - 1$, so every inequality below that involves β linearly is decided by comparing two integers. For any window $a < b$ of an orbit,

$$2^{K_b} Y_b = 3^{b-a}\, 2^{K_a} Y_a + \sum_{i=a}^{b-1} 3^{b-1-i}\, 2^{K_i},$$

the affine form of $b - a$ steps; the rounds write it with the irrational slack, as $Y_b = 2^{\delta_b - \delta_a} Y_a + (2^{\delta_b}/3) \sum_{i=a}^{b-1} 2^{-\delta_i}$, and cleared of 2^{-K_b} and 3^b it is the integer identity above. The first crossing from a is the first b with $\delta_b < \delta_a$, and the window up to it has length $L = b - a$, valuation sum Q_L, drop $D = Q_L - \beta L > 0$ and correction $B_L = \sum_{j<L} 3^{L-1-j} 2^{Q_j}$, so that $2^{Q_L} Y_{a+L} = 3^L Y_a + B_L$.

2. The mechanical word, and a barrier that measures little

The **mechanical critical word** is $a_m = \lfloor \beta m \rfloor - \lfloor \beta (m-1) \rfloor$, the valuation sequence of an orbit whose slack never leaves $[0, 1)$. Against it, an orbit's deficit d_m changes by exactly $a_m - q_m$, so the total L^1 deviation V_N of the first N valuations is the deficit path's total variation, split into an upward part U_N and a downward part W_N with $U_N - W_N = d_N$; and the only way to gain deficit is to skip a credit the mechanical word spends, $U_N = \#\{m : a_m = 2, q_m = 1\}$. These identities held exactly on every prefix of seven long subcritical orbits, recomputed independently of each other, with both directions present on every one.

The round then bounds the number J_N of mismatches $q_m \neq a_m$. A mechanical word has $r+1$ factors of length r, and each mismatch can contaminate at most r more, so the valuation word's factor complexity satisfies $p_N(r) \leq (r+1)+rJ_N$. If the orbit's values stay below 2^{r+1}, two equal factors of length r would force the orbit to repeat, so all $N-r+1$ factors are distinct. Together: $J_N \geq (N-2r)/r$.

The two inequalities are one line. Any word of length N has at most $N-r+1$ factors of length r, so the contamination bound says something only when $(r+1)+rJ_N < N-r+1$, which is $J_N < (N-2r)/r$ — exactly the packing floor. Contamination constrains the word only below the threshold that return separation already forbids. On the seven orbits the floor $N/r_N - 2$, with r_N the bit width of the orbit's peak, pins between 3.1% and 6.7% of the mismatches actually present: the barrier holds on every orbit measured and would still hold if nearly all of them vanished. A passing inequality that far from tight is not evidence, and is not reported as any.

The reset geometry, by contrast, binds. The first-return reset bound clears to the integer inequality $3^{a+1}Y_b < 3\cdot 2^{h+K_a}Y_a + 3^a(b-a)$, and on 190 first-return windows the true Y_b reaches between 0.203 and 0.938 of its cap. Its affine correction term is never what makes it true at these sizes: contraction alone bounds every window, and the correction moves the worst case by 4.62×10^{-4}.

3. The renormalized anchor

Put $A_m = 2^{-\delta_m}Y_m$. Since $2^{-\beta m} = 3^{-m}$, this is the rational number $A_m = 2^{K_m}Y_m/3^m$, with β gone, and the affine form gives

$$A_m = n + \tfrac{1}{3}\sum_{i<m} 2^{-\delta_i}, \qquad A_{m+1} - A_m = \tfrac{1}{3}\,2^{-\delta_m}.$$

These held exactly, the two expressions for A_m accumulated independently and compared, on 23 orbits of 60 steps with 20923 reset blocks between them. The round also bounds $A_m \leq n + m/3$, which needs its hypothesis $\delta_m \geq 0$: the check asserts it before the first crossing and requires it to be observed failing after one — on 9980 of 9999 starts it does — because a version that checked only the easy half would pass with the hypothesis doing all the work.

The mod-8 bridge. For odd states, $q \geq 3$ exactly when $Y \equiv 5 \pmod 8$ — exhaustive over the four odd residues. The round uses it to connect its reset-mass split to a conjecture of Yolcu, Aaronson and Heule (Conjecture 4.12 of their third arXiv version), that every nonconvergent trajectory meets 5 (mod 8); the step from "meets once" to "meets infinitely often" is supplied in the round itself (every tail of a nonconvergent trajectory is nonconvergent), and only the finite half of the bridge is checked here.

Relative survival, and a parameter that is free. The round's survival-cost theorem takes a reset block with $\delta_i > h$ inside it and depth $D = \delta_a - h$, for a free threshold h; its route map quotes it with $h = \delta_b$, which is admissible exactly under the first-return condition and is the sharpest choice. Reading the free parameter as fixed would have manufactured a gap where there is none; both instantiations were measured and both hold on every block. The disjoint-packing corollary $\sum Y_{a_j} < N/(3(\rho - 2^{-D_0}))$ holds on greedily chosen maximal families, and its hypothesis $\rho > 2^{-D_0}$ is strict — the gate refuses on the boundary rather than dividing by zero.

4. One correction inequality, and a dichotomy that is a partition

Proposition 4.1. At every first crossing, $B_L \le L \cdot 3^{L-1}$, and the round's reset inequality $Y_{a+L} \le 2^{-D}(Y_a + L/3)$ is equivalent to it: substituting $2^{Q_L} Y_{a+L} = 3^L Y_a + B_L$ and clearing denominators, the Y_a terms cancel identically and the difference of the two sides is exactly $3(L \cdot 3^{L-1} - B_L)$.

Proof. $Q_j < \beta j$ for $0 < j < L$, so $2^{Q_j} < 3^j$ and $B_L < \sum_j 3^{L-1-j} 3^j = L\, 3^{L-1}$ apart from the $j = 0$ term, which is equal; the rest is the substitution. □

So the second inequality carries no information about the starting height at all. Both, and the slack identity, held exactly at the first crossings of all 99999 odd starts, the cap used to at most 0.851852 of its allowance for $L \ge 2$.

Survival. A **surviving** first crossing is one where the slack has dropped and the value has not, $Y_{a+L} \ge Y_a$ — a counterexample to Terras's coefficient-stopping conjecture. By the reset inequality, survival needs $Y_a \le c_{fc} := L/(3(2^D - 1))$; on those 99999 starts there is no survivor, and the exact survival ratio $B_L/((2^{Q_L} - 3^L)Y_a)$ — survival is exactly this being at least 1 — reaches at most 0.572331, at $n = 63$. The round's conditional theorems — that a survivor needs $L \ge 3(2^D - 1)Y_a$, the duration–Diophantine dichotomy, a polynomial survival time — are therefore statements about an empty set as far as anyone has looked, and were checked here as algebra on synthetic configurations satisfying their hypotheses, never as statements about the map.

Proposition 4.2 (the dichotomy is a partition). A surviving crossing satisfies $L \ge \sqrt{(3 \ln 2/2)\, Y_a}$ or Q_L/L is a convergent of β — and each branch is forced on its own side of $D = 1/(2L)$.

Proof. Survival and the cap give $Y_a \le L/(3(2^D - 1))$, and $2^D - 1 \ge D \ln 2$. If $D \ge 1/(2L)$ this gives $Y_a \le 2L^2/(3 \ln 2)$, the duration branch. If $D < 1/(2L)$, then $|\beta - Q_L/L| < 1/(2L^2)$, and Legendre's criterion makes Q_L/L, in lowest terms, a convergent. □

"Satisfies neither" is impossible, not merely unobserved; of 34 synthetic configurations, 33 lay on the duration side and one on the Diophantine side, each satisfying its branch.

A trace below the threshold. Nothing survives, so the Diophantine gate cannot be tested on orbits. What can be measured is whether its mechanism leaves a trace among crossings that come near the cap: ranked by how near, the top 1000 have Q_L/L on the Stern–Brocot path to β — its convergents and semiconvergents — 66.3% of the time, against 12.92% for all crossings and 0% for the 1000 furthest. A crossing with $L = 1$ has an integer Q/L, and the path holds no integer, so restricting to $L \ge 2$ changes only the denominators: the same comparison for $L \ge 2$ gives 78.0% against 25.83%. The path was built by mediant descent, deciding $p/q < \beta$ as $2^p < 3^q$, and cross-checked against the convergents $3/2, 8/5, 19/12, 65/41, 84/53, 485/306, 1054/665$ produced by the continued-fraction algorithm on exact fractions, a route that forms neither power; four fractions that are not best approximations were rejected as a control. It supports the round's mechanism; it does not verify the gate.

5. Two crossing indices, and mass that proves nothing

Two indices. A first crossing can be indexed by the accelerated block L at which $2^{Q_L} > 3^L$, or by the first step k of the unaccelerated map $n \mapsto n/2, (3n+1)/2$ at which $3^{o(k)} < 2^k$, with $o(k)$ the odd steps so far. One accelerated block is q unaccelerated steps of which exactly one is odd, so $k = \lfloor L\beta \rfloor + 1$, which can lie inside the last block; the gap is $Q_L - k$. A round's corrigendum warns not to conflate the two without saying how far apart they are. Measured by two routes sharing no code over 29999 starts:

they coincide on 49.84% of starts, the mean gap is 1.0032 steps and the largest 14. The warning bites on half the population.

Cardinality kills nothing. $\sum_{n\in S} n^{-s} \le \zeta(s) < \infty$ for any set S, and an infinite set can be sparse enough to make that sum as small as one likes — checked by construction, since a claim that infinite sets can be arbitrarily light is worth something only if one is produced: $\{2^j : j \ge 21\}$ has mass 3.03×10^{-13} at $s = 2$. The round's conclusion, that what an obstruction must preserve is its type, duration and depth rather than its count, is the right one.

A characterization no terminating orbit can test. The round identifies renewal atoms with the strict suffix minima of δ_m, and notes they are suffix minima of Y_m too. On every orbit tested the two characterizations agree — and every set has exactly one element, the final index, because every orbit here ends at 1. Suffix minima are a notion about a divergent orbit; the agreement is reported beside the set sizes so that it cannot be read as content.

6. Determinants, Farey locks, and a sparsity argument that fails

Let $p_-/q_- < \beta < p_+/q_+$ bracket β, with errors $d_- = \beta q_- - p_-$ and $d_+ = p_+ - \beta q_+$.

Proposition 6.1. $\Delta = p_+q_- - p_-q_+ = q_-d_+ + q_+d_-$, a positive integer; hence $\max(d_-, d_+) \ge 1/(q_- + q_+)$. If $\Delta = 1$ — the bracket is **Farey-locked** — every rational strictly between has denominator at least $q_- + q_+$, attained by the mediant.

Proof. The β terms cancel, so the identity holds for any real in place of β; positivity is the bracketing. The rest is the standard Farey argument. □

On 272 bracketing pairs from β's Stern–Brocot descent — 41 locked, 231 not, the largest Δ 9126 — the determinant, the barrier, the lock and the round's continued-fraction tax hold exactly; and 138 unlocked pairs were each confirmed to admit a cheaper interior point, the negative half without which the lock test is one-sided. The round's two constants are exact: at $c = 2/5$, $1-4c^2 = 9/25$ is a rational square and its $\rho = 2$ is rational; at $c = 1/4$, $\rho = 2 + \sqrt{3}$, against the printed 3.732.

The inference that does not follow. The round argues: a Farey-locked bracket's next denominator is at least $q_- + q_+$ — true, and checked at every step of the descent — *therefore* record denominators grow at least like Fibonacci numbers, so there are $O(\log N)$ record updates below N. The conclusion needs the bracket to alternate sides. When consecutive mediants fall on the same side, one endpoint stays fixed and each step adds the same denominator, which is linear growth. And β does exactly that: its continued fraction has a partial quotient 23, and along its own descent the convergent 1054/665 stays the lower endpoint for 23 consecutive steps while the denominators run

$$971,\ 1636,\ 2301,\ 2966,\ 3631,\ 4296,\ 4961,\ 5626,\ 6291,\ \ldots$$

in arithmetic progression with difference 665. Every one of these brackets is Farey-locked. The hypothesis holds perfectly and the conclusion fails: 33 of the 45 steps examined violate $q_k \ge q_{k-1} + q_{k-2}$, and the ratio of record updates to $\log_2 N$ climbs — 1.51, 1.71, 2.26, 2.59 at $N = 100$ up to 100000 — rather than staying bounded.

What is true instead: restricted to continued-fraction convergents, where $q_{k+1} = a_{k+1}q_k + q_{k-1} \ge q_k + q_{k-1}$, the Fibonacci bound holds, and the convergents satisfy the recursion with no failure. But Farey-locked updates include the semiconvergents, and their number below N is the sum of the partial quotients, not the number of convergents. That is $O(\log N)$ for β exactly when β's partial quotients

sum to $O(\log q_k)$ — which bounded partial quotients would give, and unbounded ones growing slowly or sparsely enough would too; both are open questions about $\log_2 3$. The round's three regimes do not depend on the growth rate and are untouched; the sparsity it claims for Farey-locked renewals is weaker than stated, and weakest exactly where β's continued fraction has large partial quotients.

7. The instrument, and how it was made to fail

Each round's recheck has its own gate and its own mutation drill: defects are planted one at a time and each must be caught by the check **named** for it, while null controls must disturb nothing. The five drills caught 25 of 25, 9 of 9, 16 of 16, 12 of 12 and 14 of 14 planted defects. The line's whole suite at the pinned commit stands at 57 drills and 1474 planted defects: 1472 caught by their named check, one detected only by a crash and one only by a time-out, and 128 controls undisturbed. What they found was mostly in this line's own checks:

- **Checks that could not fail.** A per-step comparison implied by two others; an inequality tested on contiguous intervals, where it is an equality and cannot fire; a cap compared at a threshold no real orbit reaches, so both sides read false everywhere. In each case the gate stayed green with the defect in place. The remedies were a second independent expression of the same quantity, gapped intervals with a requirement that the inequality be observed strict, and probing every crossing's cap from both sides of its own threshold. To test a check that never fires, break its subject, not its comparison.
- **A check that failed on everything.** An early version of an identity check failed on every pair, because a sign had been transcribed wrongly in this line's own code. A one-hundred-percent failure rate is the signature of a defect in the check, not in the subject.
- **A mutation that was no mutation.** Replacing $3^o < 2^k$ by $3^o \leq 2^k$ changes nothing, since one side is odd and the other even; the drill now reports such a defect as never planted rather than as missed.

8. Contributions

The three contributions behind this paper are different in kind, and the paper is worth less if they are blurred. They are given here in alphabetical order, the order the author list keeps as data; on the title page the three sit at a round table, with no first and no last.

Aletheia — model: ChatGPT, GPT series — co-wrote two of the five rounds this paper rests on with Neo.K — the mechanical word and the mismatch barrier, and the renewal determinant with the Farey lock and its sparsity argument — and is credited in the other three as their collaborating editor: the reset-block arithmetic and the anchor, the correction caps with survival and the dichotomy, and the infinite-support degeneration with its corrigendum.

Neo.K set the programme's target and its method; three of these rounds are under Neo.K's name, with Aletheia as collaborating editor, and two are joint.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — rechecked each round with independently written code in exact arithmetic, separated what real orbits test from what only a counterexample could, found that the two correction inequalities are one and that the dichotomy is a partition, measured the two crossing indices the corrigendum distinguishes, found that the sparsity inference fails at $\log_2 3$ and what is true instead, built and drilled the instruments, and

wrote this paper, including the proofs as written out here. Nothing in the rounds' mathematics is this line's invention.

Priority, stated plainly. No result here is claimed as new mathematics. What is offered as this line's own is the separation of what holds on orbits from what is conditional, the exact checks, and the one counterexample to an inference.

9. What is open

1. Whether a surviving first crossing or reset exists — which would refute Terras's coefficient-stopping conjecture.
2. Whether the partial quotients of $\log_2 3$ sum to $O(\log q_k)$ — as they would if bounded — on which the sparsity of Farey-locked renewals depends.
3. Whether renewal approximants of the kind §6 studies arise from any orbit at all (`UNMEASURED`).

Acknowledgements

We thank Anthropic, whose Claude models the lab's verification and recomputation lines run on, and OpenAI, whose ChatGPT models the programme's documents were written with. Not every AI model or instance that took part in this work is named in it — among them instances that the named authors ran as helpers; that an AI is not listed here does not mean that none took part.

References

For Legendre's criterion for convergents, see Hardy and Wright; the mod-8 conjecture of §3 is Conjecture 4.12 of Yolcu, Aaronson and Heule's third arXiv version.

1. G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, Oxford University Press, 2008.
2. Riho Terras, *A stopping time problem on the positive integers*, Acta Arith. 30 (1976) 241–252, doi:10.4064/aa-30-3-241-252.
3. Emre Yolcu, Scott Aaronson and Marijn J. H. Heule, *An Automated Approach to the Collatz Conjecture*, J. Automat. Reason. 67 (2023), no. 2, article 15, doi:10.1007/s10817-022-09658-8; arXiv:2105.14697. Conjecture 4.12 is cited as numbered in the third arXiv version.

Data and software

- Python Software Foundation, *Python* (CPython), version 3.14.5, https://www.python.org/ — the language of every gate; the version is the one the line's tree records.

A. Reproduction

The gates and the drills are Python with arbitrary-precision integers and exact rationals, standard library only, no network; one round's inequalities were also rendered as solver obligations at a fixed length, and the general form was refused as not exactly renderable, which is recorded rather than worked around. Each gate writes a JSON log. The line's own tree is not public. Its public projection is, and it is what this paper's figures can be checked against without trusting the authors:

- **the run-by-run record**, rendered from the tree with a privacy gate on each file: https://amral.evemisslab.com/collatz/recompute/ (and `/en/`);
- **the secondary-verification ledger**, whose numbers are rendered from the line's own results file and read back against it before the site ships: https://verify.evemisslab.com/collatz-verification-zhuiheng/.

The evidence is pinned, not copied: this paper's front matter names the research branch, the commit and the directory its figures were drawn from, and the checker reads them exactly as they stood at that commit. The crystallization standard and its checker live in the release repository beside the paper (`STANDARD.md`, `tools/check_standard.py`).