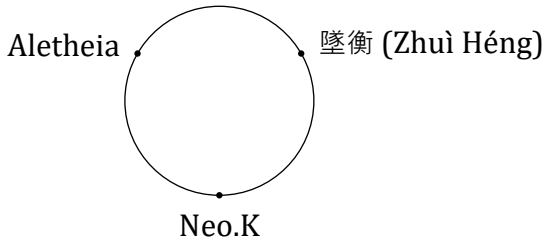
Packing deficits by congruence, and an exact certificate for a transport exponent

Four rounds of a Collatz programme, each step decided in integers or rationals:
$1/6 \to 1/9 \to 4/45 \to 1373/25856$



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. This paper concerns one quantity a Collatz programme uses to control a hypothetical divergent orbit: along a segment of the accelerated orbit that begins at a value y and ends where the orbit's slack first drops, the endpoint is $z = 2^{-D}\,\Pi$ times y, with $D > 0$ and a multiplicative correction $\Pi = \prod(1 + 1/(3Y_j))$ over the segment's states. The endpoint can lie above y only if Π is large, and four consecutive rounds of the programme bound Π by envelopes growing like a power of the segment's length L.

The exponent falls from $1/6$ to $1/9$, to $4/45$, to $1373/25856 \approx 0.0531$, and every step is decided in integers or rationals. The second step is one congruence: $3n + 1 \equiv 1 \pmod 3$, so no accelerated image is divisible by three, and the states pack at spacing three rather than two. The third is an exact identity that telescopes. The fourth is not an estimate at all but the value of a finite dual certificate, which we print in full; a reader can verify it with no computation of ours. We separate the premises carefully: the first two envelopes hold on every such segment, and were tested on tens of thousands of them; the last two are proved under the premise $z > y$, which describes the hypothetical object and which no real segment we built satisfies, so on real orbits they are consistent but untested. We give every proof the paper relies on, record what the programme's own records got wrong, and report how the instruments were made to fail.

Contents

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

1. The setting **3**

2. Three kinds of premise **4**

3. The consecutive-odd envelope, and a Gamma form that needs no Gamma **4**

4. No image is divisible by three **5**

5. Two span theorems whose premise every prefix meets **6**

6. The value-order bridge, and an exponent of $4/45$ **7**

7. A transport exponent as a certificate **8**

8. What the records got wrong **10**

9. The instrument, and how it was made to fail **10**

10. Contributions **11**

11. What is open **11**

Acknowledgements **11**

References **12**
Data and software 12

A. Reproduction **12**

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 10 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 result in this paper rests on a cited theorem.

Not claimed anywhere in this paper: the Collatz conjecture; that a divergent orbit exists or does not; that any of the four exponents is optimal; that the transport exponents α_h tend to zero as the modulus grows; that the conditional envelopes of §§6–7 were tested on a real orbit — their premise is met by none; and any exponent that passes through the programme's Diophantine input $\rho_\star = 4.1164$. The rounds also derive root-source and dense-overlap exponents from that input; the programme later withdrew it, and those exponents are left out of this paper entirely.

Inherited: nothing. Every statement below is either proved in full in the text or computed here, and every universal statement is also checked pointwise on real orbits wherever its premise is met. The four rounds' statements are the programme's; the proofs written out here are given so that the paper can be read, and checked, without the programme's documents.

round	second-layer record	what it adds	exponent of the envelope
A-U.2d.8	RUN-036	the consecutive-odd envelope, and its exact Gamma form	1/6
A-U.2d.9	RUN-037	no image divisible by three; a span theorem	1/9
A-U.2d.10	RUN-038	an exact reciprocal-flow balance; a mod-9 span theorem	4/45, conditional
A-U.2d.11	RUN-039	a residue transport identity and a dual certificate	1373/25856, conditional

1. The setting

For odd n let $q(n) = v_2(3n+1)$ and $\mathrm{Syr}(n) = (3n+1)/2^{q(n)}$, the accelerated (Syracuse) map. An orbit is $Y_0 = n$, $Y_{j+1} = \mathrm{Syr}(Y_j)$, with valuations $q_j = q(Y_j)$ and cumulative valuation $K_u = q_0 + \cdots + q_{u-1}$. Write $\beta = \log_2 3$ and call $\delta_u = u\beta - K_u$ the orbit's **slack** at step u. Since $Y_{j+1} = (3Y_j/2^{q_j})(1+1/(3Y_j))$,

$$\frac{Y_u}{Y_s} = 2^{\delta_u - \delta_s} \prod_{s \le j < u} \Big(1 + \frac{1}{3Y_j}\Big).$$

The **first crossing** of a step s is $e(s) = \min\{u > s : \delta_u < \delta_s\}$, the first time the slack falls below its value at s. A first-crossing **segment** runs from source $y = Y_s$ to endpoint $z = Y_{e(s)}$; it has length $L = e(s) - s$, states $Y_s, \dots, Y_{e(s)-1}$, total valuation $Q = K_{e(s)} - K_s$, and slack drop $D = Q - L\beta > 0$. With

$$\Pi = \prod_{s \le j < e(s)} \Big(1 + \frac{1}{3Y_j}\Big), \qquad \frac{z}{y} = \frac{3^L}{2^Q}\,\Pi = 2^{-D}\,\Pi.$$

The slack has dropped, so z can still exceed y only if $\Pi > 2^D$. The programme's route controls Π, and every round below bounds it by an envelope $E(y, L)$ that grows like a power of L/y; the exponent is what the rounds improve. A source is **post-entry** if $s \ge 1$, so that y is itself an image of the map; the **low-source** regime is $7 \le y \le L$, where the rounds state their uniform constants.

Two facts about every segment follow from the slack. For $s < u < e(s)$ we have $\delta_u > \delta_s$ (strictly, since β is irrational), so the display above gives $Y_u > Y_s$: **the source is the segment's minimum.** Its states are distinct unless the orbit closes a cycle inside the segment; every segment built here has distinct states, and the gates measure that rather than assume it. And a proper prefix of $m < L$ states is **subcritical**, $q_s + \cdots + q_{s+m-1} < \beta m$, which is the integer statement $2^Q < 3^m$ — that is what it means for the slack not yet to have dropped.

2. Three kinds of premise

Every theorem in these rounds is a statement about segments, and its premise decides what a test on real orbits can show. The second-layer gates measured each premise separately, printed its denominator, and applied a theorem only where its premise held. Three kinds turned up.

Met by every segment, by construction. That the source is the minimum and the states distinct (§1) held on all 32207 segments of the first round's sweep and all 23689 of the second's, with no segment failing either. Subcriticality held on all 9517 proper prefixes the second round examined and all 27345 the third did. Theorems with these premises were genuinely tested.

Met by no real segment. The premise $z > y$ — the endpoint above the source, although the slack has dropped — is the defining property of the programme's hypothetical object, a segment of an orbit that never comes down. No real first-crossing segment we built satisfies it: none of 68311 in the third round's sweep and none of 38101 in the fourth's. The depth caps of the first two rounds need its analogue for nested chains of segments, $z_1 > y_r$, the outermost endpoint above the innermost source; it was met by none of 24010 chains and none of 15832. A theorem with such a premise applied to real orbits is applied to nothing.

Not a premise at all. The identities of §§3, 6 and 7 hold on every segment unconditionally.

The middle kind is the one that is easy to get wrong, because "zero violations" and "nothing tested" produce identical logs. Applied without its premise, the first round's depth cap flagged thousands of chains as violating a theorem that holds; the honest number is that it was applied to none. So each gate also carries a planted defect that deletes its premise gate and must turn the run red — which is what shows the gate is load-bearing rather than merely exclusive. At the fourth round the same defect produced no violation at all: the certified bound has so much slack on real segments that its conclusion survives losing its hypothesis. That is a fact about the bound, stated rather than engineered around, and the defect was made decisive by also tightening the bound's constant.

3. The consecutive-odd envelope, and a Gamma form that needs no Gamma

Proposition 3.1 (first round). For every segment with distinct states,

$$\Pi \le E_{\mathrm{odd}}(y, L) := \prod_{k=0}^{L-1} \Big(1 + \frac{1}{3(y+2k)}\Big).$$

Proof. The states are distinct odd integers at least y (§1), so the k-th smallest is at least $y + 2k$, and $1 + 1/(3x)$ is decreasing. □

It held on all 32207 segments of the sweep, built from 1249 orbits, with segments up to $L = 51$ and sources up to 425645; no sorted state fell below $y + 2k$. Each factor is $(3y + 6k + 1)/(3y + 6k) = (y/2 + k + 1/6)/(y/2 + k)$, so

$$E_{\text{odd}}(y, L) = \frac{(y/2 + 1/6)_L}{(y/2)_L} = \frac{\Gamma(L + y/2 + 1/6)\,\Gamma(y/2)}{\Gamma(y/2 + 1/6)\,\Gamma(L + y/2)},$$

a Pochhammer ratio, which grows like $L^{1/6}$ for fixed y. The round states this Gamma form as a theorem and its checker verifies it numerically, reporting a worst error of 3.1×10^{-10}. There is nothing to verify numerically: at integer L the functional equation turns the Gamma quotient into the product it was built from, and the two sides are one rational number. Checked as rationals, the disagreement is zero on every segment and on all 56 designed pairs up to $L = 400$. An independent numeric check with the library log-gamma function agreed too, but only after its tolerance was set by the cancellation it actually suffers: $\ln E_{\text{odd}}$ is a difference of four large log-gammas that nearly cancel, and a relative-error test under that much cancellation had reported twelve correct pairs as failures.

4. No image is divisible by three

The first round listed as open that a real orbit should beat the consecutive-odd envelope strictly, with no quantitative statement available. The second round closes it in one line.

Proposition 4.1. For odd n, $\text{Syr}(n)$ is coprime to 6. Hence every post-entry state is 1 or $5 \pmod 6$.

Proof. $3n + 1 \equiv 1 \pmod 3$, and dividing by a power of two introduces no factor of three; the image is odd by construction. □

Corollary 4.2. A post-entry source with $L \geq 2$ is 7 or $11 \pmod{12}$.

Proof. If $q_s \geq 2$ then $\delta_{s+1} = \delta_s + \beta - q_s < \delta_s$ and the segment ends after one step; so $L \geq 2$ forces $q_s = 1$, which on odd integers is $y \equiv 3 \pmod 4$. Combine with Proposition 4.1. □

Checked directly: of 20000 odd integers mapped, no image was divisible by three; of 24688 post-entry states, none lay outside $\{1, 5\} \pmod 6$; of 11914 post-entry sources with $L \geq 2$, none lay outside $\{7, 11\} \pmod{12}$. The premise $L \geq 2$ is load-bearing: every one of the 11775 sources with $L = 1$ lies outside that pair, as the proof says it must, and a first version of the check that forgot the premise reported all 11775 as violations of a corollary that holds.

So the states are packed into the integers coprime to 6 — mean spacing three, not two. Let $a_0(y) < a_1(y) < \cdots$ be the integers at least y coprime to 6; for post-entry y, $a_{2m} = y + 6m$ and $a_{2m+1} = y + 6m + c$ with $c = 4$ if $y \equiv 1$ and $c = 2$ if $y \equiv 5 \pmod 6$.

Proposition 4.3. For every segment with distinct states and a post-entry source, $\Pi \leq E_6(y, L) := \prod_{k<L}(1 + 1/(3a_k(y)))$, and $E_6(y, L) = \Theta_y(L^{1/9})$.

Proof. As for Proposition 3.1, with the k-th smallest state at least $a_k(y)$. Splitting the product by the parity of k gives two Pochhammer ratios $(b + \frac{1}{18})_n/(b)_n$ with $b = y/6$ or $(y + c)/6$ and $n \approx L/2$, each growing like $n^{1/18}$; whence the exponent $\frac{1}{9}$. □

It held on all 23689 segments; the two-progression Gamma form agreed with the product exactly, on every segment and on all 42 designed pairs. The deficit against the odd envelope is a power, $E_6/E_{\text{odd}} = \Theta(L^{-1/18})$. Table 1 shows the finite-L exponents $\ln E / \ln L$ moving toward their limits; these are trends, not the theorem.

y	L	E_{odd}, toward 1/6	E_6, toward 1/9	E_{RF}, toward 4/45
7	200	0.13160	0.09645	0.12640
7	12800	0.14675	0.10278	0.10986
11	200	0.11626	0.08777	0.10537
11	12800	0.13799	0.09786	0.09805

Table 1: $\ln E / \ln L$ for the three envelopes (E_{RF} is defined in §6). At every length shown the third envelope is still above the second; its gain is asymptotic.

The same round's depth cap for nested chains needs $z_1 > y_r$ and was applied to none of its 15832 chains. Its combinatorial half does not need orbits and was enumerated: over 499149 sets of $\{7, 11\} \pmod{12}$ anchors at depths 2 to 6, no span fell below the proof's allowance $6(r-1)-2$, and the phase rule that closes the remaining gap held on every tight set.

5. Two span theorems whose premise every prefix meets

Both rounds that follow also prove a statement about the **span** W of a proper prefix — the difference between its largest and smallest state, plus one — whose only premise is subcriticality, and so is met by every prefix of every segment. Both proofs are short.

Lemma 5.1. On odd integers, $q(n) = k$ exactly when n lies in one residue class modulo 2^{k+1}, namely $n \equiv 3^{-1}(2^k - 1)$. Among W consecutive integers, at most $W/(3 \cdot 2^k) + 2$ are coprime to 3 and lie in that class.

Proof. $v_2(3n+1) = k$ iff $3n + 1 \equiv 2^k \pmod{2^{k+1}}$. Excluding multiples of three leaves two classes modulo $3 \cdot 2^{k+1}$, and each meets W consecutive integers at most $W/(3 \cdot 2^{k+1}) + 1$ times. □

Proposition 5.2 (second round). A subcritical prefix of m distinct post-entry states with span W satisfies $W > \frac{24}{17}\big((4-\beta)m - 12\big)$.

Proof. Let N_k count the prefix's states with $q = k$. Each state contributes at least $\min(q, 4)$, so $\sum q \geq 4m - 3N_1 - 2N_2 - N_3$, and by Lemma 5.1 $3N_1 + 2N_2 + N_3 \leq 3(\frac{W}{6} + 2) + 2(\frac{W}{12} + 2) + (\frac{W}{24} + 2) = \frac{17}{24}W + 12$. Subcriticality gives $\beta m > \sum q \geq 4m - \frac{17}{24}W - 12$. □

The coefficient $17/24$ is where the sieve does its work: without the factor three in Lemma 5.1 every class would be half again as dense, and so would the coefficient. The statement held on all 9517 subcritical prefixes, up to $m = 50$; on every enumerated window the per-class count even met the tighter bound $W/(3 \cdot 2^k) + 1$.

The third round adds a cost read modulo 9. If $Y_{j+1} = t$ is the image of Y_j, then $2^{q_j} t = 3Y_j + 1 \equiv 4$ or $7 \pmod 9$, since Y_j is coprime to 3; as 2 has order 6 modulo 9, the class of t fixes a least admissible valuation, its **cost**: $1, 2, 2, 3, 4, 1$ for $t \equiv 2, 1, 4, 5, 7, 8 \pmod 9$ in turn — that is, cost one for $t \equiv 2, 8$, cost two for $t \equiv 1, 4$. The six entries were rederived from the congruence, not transcribed; no edge of 68311 fell below its cost, and no image fell outside $\{4, 7\}$ after multiplying back.

Proposition 5.3 (third round). If the m images of a subcritical prefix are distinct with span W, then $W > 3(3-\beta)m - 18$.

Proof. Odd integers in two classes modulo 9 are two classes modulo 18, at most $W/9 + 2$ of them among W consecutive integers. So $\sum q \geq 3m - 2(\frac{W}{9} + 2) - (\frac{W}{9} + 2) = 3m - \frac{W}{3} - 6$, and subcriticality gives $\beta m > 3m - W/3 - 6$. □

The valuation floor and the span bound held on all 27345 subcritical prefixes. Neither span theorem is promoted to an exponent for Π: a span controls the diameter of the realized states, not where their reciprocal mass sits, and the rounds forbid the promotion in so many words.

6. The value-order bridge, and an exponent of $4/45$

The step from a diameter to a product needs value-order information, and the third round finds it in the edge equation itself.

Proposition 6.1 (identity). For every edge,

$$\frac{1}{Y_j} - \frac{1}{Y_{j+1}} = \frac{3 - 2^{q_j}}{3Y_j} + \frac{1}{3Y_jY_{j+1}},$$

and summed over a segment, $\sum_j (2^{q_j} - 3)/Y_j = -3/y + 3/z + C$ with $C = \sum_j 1/(Y_jY_{j+1})$.

Proof. Multiply by $3Y_jY_{j+1}$ and use $2^{q_j}Y_{j+1} = 3Y_j + 1$; then sum, and the left side telescopes. □

It held exactly on all 70810 edges of 2499 orbits and on all 68311 segments. The identity couples each edge's valuation to its source's reciprocal weight. Write S_k for the reciprocal mass $\sum 1/Y_j$ over the states with $q_j = k$, and S for the total. The left side is $-S_1 + S_2 + 5S_3 + 13\,S_4 + \cdots$, so **when** $z > y$ it is less than C, and

$$S_2 + 5\sum_{k\geq 3} S_k < S_1 + C.$$

Expensive valuations must be paid for out of the cheap class $q = 1$. That this balance is equivalent to $z > y$ is algebra, and it held on all 68311 segments; the balance itself was applied to none, because none has $z > y$.

Proposition 6.2 (third round). For a segment with distinct states, $z > y$ and a post-entry source $y \geq 7$,

$$S < \tfrac{1}{5}\ln\Big(1+\tfrac{6L}{y}\Big) + \tfrac{1}{15}\ln\Big(1+\tfrac{12\,L}{y}\Big) + \frac{289}{70\,y}, \qquad \Pi < E_{\mathrm{RF}}(y,L) := e^{289/(210\,y)}\Big(1+\tfrac{6L}{y}\Big)^{1/15}\Big(1+\tfrac{12\,L}{y}\Big)^{1/45}.$$

Proof. States with $q = 1$ are 7 or 11 (mod 12), those with $q = 2$ are 1 or 17 (mod 24) (Lemma 5.1 with the sieve). At most L distinct members of one progression of step d, all at least y, have reciprocal mass at most $1/y + d^{-1}\ln(1 + dL/y)$; splitting between two progressions and using concavity gives $S_1 \leq 2/y + \frac{1}{6}\ln(1 + 6L/y)$ and $S_2 \leq 2/y + \frac{1}{12}\ln(1 + 12\,L/y)$. With $z > y$ the endpoint is at least y too, so the $L + 1$ states are distinct odd integers at least y, and $1/(ab) \leq (a^{-2} + b^{-2})/2$ gives $C \leq \sum_{n\geq y,\ n\ \mathrm{odd}} n^{-2} \leq 1/y^2 + 1/(2y) \leq 9/(14\,y)$. The balance gives $\sum_{k\geq 3} S_k < (S_1 + C - S_2)/5$, hence $S < \frac{6}{5}S_1 + \frac{4}{5}S_2 + \frac{1}{5}C$; substitute, and note $\frac{6}{5}\cdot 2 + \frac{4}{5}\cdot 2 + \frac{1}{5}\cdot\frac{9}{14} = \frac{289}{70}$. Finally $\ln\Pi \leq S/3$. □

So the exponent is $\frac{1}{15} + \frac{1}{45} = \frac{4}{45}$, a gain of $1/45$ over $1/9$; in the low-source regime, $\Pi < C_{10}(L/y)^{4/45}$ with $C_{10} = e^{289/1470}\,7^{1/15}\,13^{1/45}$.

The two capacity bounds need nothing beyond §1's structure and held on all 68311 segments — 33939 of which have no $q = 1$ edge at all. The total-mass bound and the product envelope are conditional, and we record a correction to the second layer's account of this round here: its record grouped

the total-mass bound with the unconditional capacities, but the proof above passes through the balance and through the cross-term bound, and both need $z > y$. Evaluated on the real segments anyway, the total-mass bound was never exceeded on any of the 68311, and the uniform form never on any of the 241 low-source segments. That is consistent with the proposition and is not a test of it.

The round also limits itself. It builds a relaxation that satisfies every static one-point valuation constraint yet keeps the exponent 1/9, showing that the gain cannot come from valuation counts alone. Its two closed forms, a density $D(t) = 1/(3(2-t))$ and an average valuation tending to $2/(2-t)$, eliminate to a relation between the round's own reported figures, $\overline{q} \to 6|S_X|/X$. At the three reported sizes $X = 30000$, 100000 and 300000 — counts 7137, 23802 and 71423 — the gap is 2.01×10^{-3}, 7.19×10^{-4} and 1.81×10^{-4}, shrinking as the $o(X)$ in $|S_X| = D(t)X + o(X)$ predicts, and the recovered t converges to 0.5999. A first attempt at this check enumerated a different set and printed its average beside theirs, which reads like a reproduction and is not one.

7. A transport exponent as a certificate

The fourth round lifts the balance from one class to residues modulo $M = 3^h$. Every state is a unit modulo M (Proposition 4.1), and an edge from a state $\equiv r$ with valuation k lands in the class $T_h(r,k) = (3r+1)2^{-k} \bmod M$, again a unit.

Proposition 7.1 (transport identity). For every segment and every unit class b modulo M,

$$3S_b - \sum_{T_h(r,k)=b} 2^k S_{r,k} = \frac{3}{y}[y \equiv b] - \frac{3}{z}[z \equiv b] - C_b,$$

where S_b and $S_{r,k}$ are the reciprocal masses of the states $\equiv b$, and of those $\equiv r$ with valuation k, and C_b sums $1/(Y_j Y_{j+1})$ over the edges landing in b.

Proof. $3/Y_{j+1} = 2^{q_j}/Y_j - 1/(Y_j Y_{j+1})$ for every edge; sum over the edges landing in b, and compare the targets with the sources, which differ only at the two ends. □

It held exactly on all 685818 residue identities modulo 27 checked, one per unit class per segment, over 38101 segments of 1499 orbits.

A **level-h certificate** is a positive potential a_r on the units modulo 3^h and non-negative multipliers $\mu_{r,k}$, finitely many non-zero, with

$$-3a_r + 2^k a_{T_h(r,k)} + \mu_{r,k} \geq 1 \qquad \text{for every unit } r \text{ and every } k \geq 1.$$

Proposition 7.2 (fourth round). Given a level-h certificate, every segment with distinct states, $z > y$ and a post-entry source $y \geq 7$ satisfies

$$S \leq \frac{51\, a_{\max}}{14\, y} + \sum_{r,k} \mu_{r,k}\Big(\frac{1}{y} + \frac{1}{D_k} \ln\Big(1 + \frac{D_k L}{y}\Big)\Big), \qquad D_k = 3^h 2^{k+1},$$

and therefore $\Pi = O_y(L^{\alpha_h})$ with

$$\alpha_h = \frac{1}{3} \sum_{r,k} \frac{\mu_{r,k}}{3^h\, 2^{k+1}}.$$

Proof. Weight the identity of Proposition 7.1 by a_b and sum over b: $\sum_{r,k}(3a_r - 2^k a_{T_h(r,k)})S_{r,k} = 3a_{[y]}/y - 3a_{[z]}/z - \sum_b a_b C_b$. The certificate gives $S_{r,k} \le (2^k a_{T_h(r,k)} - 3a_r + \mu_{r,k})S_{r,k}$; summing, $S \le 3a_{[z]}/z - 3a_{[y]}/y + \sum_b a_b C_b + \sum \mu_{r,k} S_{r,k}$. With $z > y$ the boundary is at most $3a_{\max}/y$ and the cross terms at most $a_{\max} \cdot 9/(14\,y)$, as in §6. By Lemma 5.1, the states $\equiv r \pmod{3^h}$ with valuation k form one class modulo D_k, so $S_{r,k} \le 1/y + D_k^{-1} \ln(1 + D_k L/y)$. Finally $\ln \Pi \le S/3$. □

The exponent is a finite rational computed from the certificate, and whether a certificate is valid is a finite check: past the smallest K with $2^K a_{\min} - 3a_{\max} \ge 1$ the inequality holds from the transport term alone, so only $k < K$ needs checking. **Nothing in either half needs a reference computation of ours.**

The level-1 certificate fits in a line. Modulo 3 the units are 1 and 2, and $T_1(r,k)$ is 2 for odd k and 1 for even k. Take

$$a_1 = \frac{1}{10}, \quad a_2 = \frac{1}{5}, \qquad \mu_{1,1} = \mu_{1,2} = \frac{9}{10}, \quad \mu_{2,1} = \mu_{2,2} = \frac{6}{5}.$$

For $r = 1$ the inequality reads $-\frac{3}{10} + \frac{2}{5} + \frac{9}{10} = 1$ at $k = 1$ and at $k = 2$, and $\frac{13}{10}$ at $k = 3, 4$; for $r = 2$ it reads 1 at $k = 1, 2, 3, 4$. The tail starts at $K = 4$, since $2^4 \cdot \frac{1}{10} - \frac{3}{5} = 1$. Then $\alpha_1 = \frac{1}{3}\left(\frac{9/10}{12} + \frac{9/10}{24} + \frac{6/5}{12} + \frac{6/5}{24}\right) = \frac{7}{80}$.

The round ships certificates at three levels. All three were verified exactly: every potential positive, every multiplier non-negative, every transition inside the unit group, the inequality checked for every k up to five past the tail, and each level's declared tail exactly the computed K.

level	modulus	units	non-zero μ	tail K	inequalities	violations	α_h
1	3	2	4	4	18	0	7/80
2	9	6	8	5	60	0	99/1472
3	27	18	18	7	216	0	1373/25856

Table 2: the three certificates; 294 inequalities, no violation.

The level-3 certificate, quoted verbatim from the round's certificate file, is:

r	a_r	r	a_r	r	a_r
1	457/808	10	1121/1616	19	85/202
2	457/404	11	85/101	20	1121/808
4	239/404	13	919/808	22	89/101
5	46/101	14	4979/12928	23	85/404
7	4979/25856	16	23/101	25	85/808
8	239/202	17	178/101	26	919/404

with non-zero multipliers $\mu_{1,k} = 351/808$, $\mu_{2,k} = \mu_{22,k} = 351/404$, $\mu_{17,k} = 351/202$, $\mu_{10,k} = 1323/1616$, $\mu_{13,k} = \mu_{20,k} = 1323/808$, $\mu_{26,k} = 1323/404$, each for $k = 1, 2$, and $\mu_{13,3} = \mu_{13,4} = 621/808$. It gives $\alpha_3 = 1373/25856 \approx 0.0531$, and the gain over $4/45$ is exactly $41639/1163520$. The exponents decrease with the level; the round's floating values for deeper moduli continue the decrease, but it labels them diagnostics and does not claim they tend to zero, and neither do we.

On real segments the uniform level-3 envelope was never exceeded on any of the 141 low-source segments, and Proposition 7.2 itself was applied to none: as in §6, its premise is $z > y$.

8. What the records got wrong

None of the following changes a result; each is the kind of thing a reader using the programme's machine-readable files would otherwise inherit.

- **An exact identity checked as an approximation.** The Gamma forms of §§3–4 were verified by the programme's checkers numerically, with reported errors of 3.1×10^{-10} and 2.0×10^{-13}, where the identity is exact and the error is zero.
- **Rationals published as their floating evaluations.** Constants with exact rational values were published as the float64 result of evaluating their formulas, and constants derived from them were computed from the already-rounded parent rather than from their closed forms, so the last-bit error passed down a chain: in the second round $C_9 = C_6/6$ and $c_9 = C_9^{-9/8}$ each sit one unit in the last place from the correctly rounded value, and the error reproduces exactly by redoing the arithmetic in doubles on the published parent. The pattern recurs in all four rounds. In the third, the span coefficient $24(4-\beta)/17$ appeared with two different values one unit apart in two files of one bundle; the value carried into the next round was the correct one.
- **A ledger without its open problems.** The fourth round's paper ends its ledger with six entries, four of them open problems; the machine-readable ledger keeps only the diagnostics among them and lists no open problem at all, so the Collatz conjecture itself is absent from the file a downstream reader would consult first. The same round ships the script that generated every other file in it, and that script is the one file no digest in the bundle pins.

9. 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 four drills caught 25 of 25, 31 of 31, 35 of 35 and 33 of 33 planted defects, each with its two controls undisturbed. 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. Exact brackets for $\ln 2$, $\log_2 3$ and the roots are certified once and imported; each gate also runs self-checks on them, eight to ten per round, with defects aimed at the brackets themselves.

What the drills and the checks found was mostly in this line's own instruments, and four instances are worth stating because each would have published a wrong sentence:

- **A rendering measured instead of a number.** Comparing a published double against a rational bracket by printing the bracket to a fixed number of places measured the printing: two constants came back tens and tens of thousands of units in the last place off, purely because the printed places were coarser than a double. Rounding is monotone, so the bracket decides by itself whether one double is nearest, and both differences fell to at most one unit.
- **A premise dropped.** Theorems applied without their premises flagged thousands of chains in the first round, all 11775 single-step sources in the second, and hundreds of segments in the third, each against a statement that holds. Each premise is now measured and printed, and a drill defect deletes it and requires the run to go red.

- **A check that could never fire.** A counter for an envelope falling below one could only read zero, since every factor exceeds one; the drill's malformed-defect class exposed it, and it was removed rather than kept as reassurance.
- **Wide is not wrong.** A defect truncating a series to two terms was caught, but not by the containment check named for it: a two-term bracket is wide, not false, and still contains the true value. What it breaks is decidability. Naming the wrong check would have scored a working guard as a miss.

10. 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 — wrote, with Neo.K, the four rounds this paper rests on: the multiplicative identity and the consecutive-odd envelope; the sieve, the admissible positions and the span theorem; the reciprocal-flow balance, the mod-9 cost, the span bound and the static countermodel; and the transport identity, the dual certificates and the exponent they give.

Neo.K set the programme's target and its method, and wrote its documents with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — rechecked each round with independently written code in exact arithmetic, measured every premise separately and so found which statements real orbits can test, found that the Gamma forms are exact, eliminated the countermodel's closed forms into a checkable relation, verified the certificates and computed their tails, quoted them into an evidence file so that they could be printed, corrected its own record of the third round's conditional statements, 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. The envelopes, the identities, the certificates and the exponents are the programme's. What is offered as this line's own is the separation of premises, the exact checks, and the record of what held where.

11. What is open

1. Whether $\alpha_h \to 0$ as $h \to \infty$, which the round's diagnostics suggest and do not prove.
2. Whether better certificates exist at the same modulus: what was verified is that the shipped ones are **valid**, not that they are optimal.
3. Whether any real or hypothetical segment meets $z > y$; on real orbits none of ours does, so the conditional envelopes of §§6–7 remain `UNMEASURED` there.
4. The exponents that pass through $\rho_\star$. A later round of this line found that the theorem behind $\rho_\star$ does bound β, by its published review, at an exponent that does not quite reach $\rho_\star$, and re-evaluated those exponents there; they are reported with the rounds that carry them, not here.

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

The theorem behind $\rho_\star$ in open item 4 is Wu and Wang's, and its published review is zbMATH Open's.

1. Qiang Wu and Lihong Wang, *On the irrationality measure of log 3*, J. Number Theory 142 (2014) 264–273, doi:10.1016/j.jnt.2014.03.007; reviewed in zbMATH Open as Zbl 1303.11081.

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, the drills and the extraction of the certificates are Python with arbitrary-precision integers and exact rationals, standard library only, no network. Each gate writes a JSON log; the certificates printed in §7 are quoted from the round's file into a separate evidence file, after its digest is matched against the bundle's own manifest and the gate's verification has run again on what is quoted. 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`).