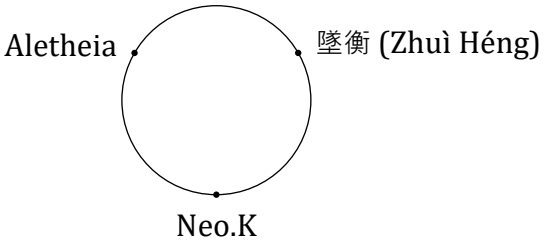
Renewal sparsity in the slack of the Collatz map: laminar crossings, a transport hierarchy, a dense branch closed and records pinned by slack — on a Diophantine input read from the review of the paper it cites

Seven rounds of a Collatz programme checked in exact arithmetic on real orbits where their objects occur, and as algebra where they do not; the one external exponent the rounds share, withdrawn by the programme, restored at the value its source's published review supports, with every exponent derived from it re-evaluated there



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. A research programme studies a hypothetical divergent orbit through its slack, $\delta_m = m \log_2 3 - K_m$, where K_m is the sum of the first m valuations of the Syracuse map, and through the renewal process of its suffix minima. Seven of its rounds build a theory of how sparse those renewals must be, and every quantitative step in them passes through one external input, a Diophantine exponent $\rho^\star = 4.1164$ for $\log_2 3$. The programme later withdrew that input, judging the cited theorem to concern $\log 3$ rather than $\log_2 3$. Its published review shows the cited result is a lower bound for linear forms in 1, $\log 2$, $\log 3$, which bounds $\log_2 3$ directly; but the review prints the exponent as 4.117, rounded up, which does not reach 4.1164. This paper restores the input at $4117/1000$, re-evaluates as exact rationals every exponent the rounds derived from it — each first reproduced exactly at $\rho^\star$ — and finds each move the weaker way with every side condition intact. Around that, the structural core is checked where it is testable: the first-crossing intervals of any orbit are laminar and an annulus identity holds as an exact integer pair; a hierarchy of block exponents goes to zero while its constants explode, so that the new bound overtakes the old only at ratios no finite orbit approaches; the dense branch of the survivor space is empty; every suffix minimum has a first step of valuation 1; a record bound needing no premise is attained; and the programme's correction bank, written in its own coordinates, is an identity in integers.

Contents

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

1. The setting **3**

2. The one external input, and what its source says **4**

3. Crossings are laminar, and the annulus identity is an integer pair **5**

4. A transport hierarchy whose exponent goes to zero **5**

5. The dense branch is empty **6**

6. Sparse support, and the price of anything denser **6**

7. Records pinned by slack **6**

8. The correction bank is an integer identity **7**

9. What real orbits can and cannot test here **7**

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

11. Contributions **7**

12. What is open **8**

Acknowledgements **8**

References **8**

Data and software 8

A. Reproduction **9**

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 11 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

label	meaning
`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. Every exponent of §§2–8 that depends on the Diophantine input is of the second kind: it holds by the cited Wu–Wang theorem, read through its published review (§2), and is written that way.

Not claimed anywhere in this paper: the Collatz conjecture, or that a divergent orbit exists or does not; the rounds' asymptotic theorems — every $o(1)$, every "for N large" — which are checked in their finite content and their exponent algebra only; that any exponent is optimal; the exact exponent of the Wu–Wang theorem, which only the paper states; and that the theorems about the hypothetical divergent branch have been tested on real orbits where their objects do not occur, which in several cases they do not.

Inherited: the asymptotic halves of the rounds' statements; the programme's reduction of a divergent orbit to its slack and renewal process; and the constant C_ε of the block-hierarchy product theorem, which later rounds consume and which is never evaluated.

Where the other labels sit: the integer identities, counts and populations of §§3–8 are `COMPUTED_HERE`; the closed forms of the exponent family, and the transfer of the review's bound to $\log_2 3$, are `DERIVED_FROM_VERIFIED_INPUTS`; the exponents at the restored input are `PASS by cited theorem`; the programme's withdrawal of the input is `FAILED` as a reason (§2); and the paper's exact exponent, and whether its "sufficiently large" is effective, are `UNMEASURED`.

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+1}, partial sums K_m, $\beta = \log_2 3$, 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 is linear in β is decided by comparing two integers, and none needs a logarithm.

The rounds read a hypothetical divergent orbit through two objects. The **first crossing** from s is $e(s) = \min\{u > s : \delta_u < \delta_s\}$: the first time the orbit has spent more halvings than the slope β allows. A **suffix minimum** is a state below every later state; the suffix minima of an orbit form its renewal process, split by the programme into A-renewals, which have no crossing, and B-renewals, which do. The seven rounds — A-U.2d.3, 4, 12, 13, 14, 15 and 16 in the programme's numbering — ask how sparse the renewals of a divergent orbit would have to be, and each answer below is either tested on real orbits, where the object occurs, or checked as algebra, where it does not.

2. The one external input, and what its source says

The input. From A-U.2d.3 on, the rounds take one external exponent: $\rho^\star = 4.1164$, with $\|q\beta\| \gtrsim q^{-\rho}$, and derive everything else from $\theta = 1/(\rho + 1)$. Because 4.1164 is a terminating decimal, every derived exponent is an exact rational, and this line checked them as such: $\theta^\star = 2500/12791$, $\sigma^\star = 1/(1 + \theta^\star) = 12791/15291$, and so on (`COMPUTED_HERE`). Two closed forms the rounds never state make the family one-parameter: the backbone exponent is $1/(1+\theta)$, and the support exponent $\kappa_{13} = (\rho + 1)/(\rho + 1 + \theta)$ is $1/(1 + \theta^2)$.

The withdrawal, and the review. The programme's last round withdrew $\rho^\star$, on its own audit that the cited theorem of Wu and Wang concerns $\log 3$, not β. This line had read a survey attributing the bound to $\log 2/\log 3$, whose irrationality measure is β's. The paper itself is behind a paywall. Its zbMATH review is not, and says the paper's bound for $\log 3$ follows from a more general result: a lower bound $H^{-4.117}$ for $|p + q_1 \log 2 + q_2 \log 3|$, with $H = \max(|q_1|, |q_2|)$ sufficiently large. Taking $p = 0, q_1 = -a, q_2 = b$, the form is $\log 2 \cdot |b\beta - a|$; since $3^5 < 2^8$ the nearest integer a to $b\beta$ is at most $2b$, so $|b\beta - a| > (2b)^{-4.117}/\log 2$ — an exponent for β (`DERIVED_FROM_VERIFIED_INPUTS`, each step an integer comparison). The withdrawal's reason does not hold for the general result (`FAILED` as a reason). But the review prints its exponent as 4.117, a bound rounded up, and $4.11633052 < 4.1164 < 4.117$: the survey's figure supports the programme's value, the review's does not reach it.

The input restored, at the value the review supports. The exponents below are therefore evaluated at $\rho_c = 4117/1000$, where $\theta_c = 1000/5117$ and $\sigma_c = 5117/6117$. Every one of the 29 archived exponents of the rounds, and the dense source floor at 10 levels, was first reproduced exactly from its formula at $\rho^\star$ — no disagreement — and then evaluated at ρ_c. A larger exponent is the weaker hypothesis, so every exponent of a lower bound must fall and every exponent of an upper bound must rise, and each one whose role is stated does. The side conditions the arguments need all hold at ρ_c: $\rho > 3$, $\rho < 5$, and the survey's $4.11633052 \le \rho < 7.616$, Rhin's older bound (`PASS by cited theorem`).

exponent	role	at $\rho^\star$	at ρ_c
θ, the survival duration $L(y) \gtrsim y^\theta$	lower bound	2500/12791	1000/5117
σ, the sparsity $R \le (N/c)^\sigma$	upper bound	12791/15291	5117/6117
$\rho/(\rho + 1)$, congestion and active backlog	upper bound	10291/12791	4117/5117
κ_{13}, completed support $M \lesssim N^{\kappa_{13}}$	upper bound	163609681/169859681	26183689/27183689
$\chi(\sigma) = (5\sigma - 4)/3$, partial-quotient pressure	—	2791/45873	1117/18351

Second-order coefficients, and a probe that can reject a neighbour. The first of the rounds quotes two second-order coefficients of asymptotic expansions. A wrong second coefficient is quiet — the leading term still dominates — so each was checked by a scaled residual that must vanish for the right value and not for a wrong one: at 10^{12} the residuals are 1.73126×10^{-7} and 3.00283×10^{-8}, while a value 1% off converges instead to -0.00706435 and -0.00499765 (`COMPUTED_HERE`). Its continued

fraction of $\log_2 3$ was recomputed with no logarithm at all: every partial quotient is a comparison of a power of 2 with a power of 3, giving 16 terms certified up to denominator 10781274.

3. Crossings are laminar, and the annulus identity is an integer pair

The fourth round's core is not about a divergent orbit at all. Its Theorem 3.1 — $e(s)$ is the next smaller element of the slack sequence — is the crossing condition rearranged, and its laminarity theorem holds because next-smaller-element intervals are laminar in any sequence. So the core applies to real orbits, and was tested on four (COMPUTED_HERE):

	count
intervals tested for Theorem 3.1, violations	394, **0**
interval pairs that properly cross	**0**
nested pairs / disjoint pairs, so both branches are inhabited	2757 / 23429
nested edges on which the annulus identity was evaluated	229

The annulus identity $A + D = D' + E$ is a relation among quantities of the form $c\beta + k$, so its residual is not a small number but the pair $(0, 0)$, on every edge. The programme's own checker evaluates the same identity in floating point and reports a residual of $2.3092638912203256 \times 10^{-14}$. The congestion envelope this round derives from $\rho^\star$, $\omega_N \lesssim N^{\rho/(\rho+1)}$, is at ρ_c the exponent $4117/5117$.

4. A transport hierarchy whose exponent goes to zero

The twelfth round replaces a single transport step by blocks of m consecutive steps and counts the exact words a shrinking block can use. Every symbol of its headline is an integer or a ratio of integers: $q_m = \lfloor \beta m \rfloor$ is the bit length of 3^m less one, $C_m^- = \sum_{Q=m}^{q_m} \binom{Q-1}{m-1} / (3 \cdot 2^Q)$, $\gamma_m = 2^{q_m+1}/3^m - 1$, and the block exponent $\hat{\alpha}_m = \frac{1}{3}(1 + 1/\gamma_m)\, C_m^-$. The six published levels and the whole record set to $m = 150$ reproduce exactly; the sequence is not monotone — it set a new low at only 57 of the 150 levels, and rose from the level before at 62 steps — and its record structure is the claim, and holds. The generating identity was checked by convolving the geometric law directly rather than through its closed form, and both halves of the closure hold: the Chernoff decay, tightest at $m = 2$ where the true value is 0.558 of the bound, and the Diophantine floor on the phase gap (COMPUTED_HERE).

Three findings bound what the hierarchy says. **A premise stated more weakly than it is used**: §1 exempts a segment's endpoint from the bound "every state at least y", while §4's sums run to the endpoint; under §1's reading the closure fails 221 times out of 999 on real orbits, because the exempted closing block can carry an unbounded valuation, and under the reading §4 uses it holds everywhere. **A premise met once**: the hypotheses of §§7–8 hold for exactly one of the 7500 segments built from 19998 sources scanned, the one from $y = 31$. **An exponent going to zero with constants going to infinity**: $\hat{\alpha}_{12} = 188368/4654215$ is below the previous round's $1373/25856$, but with the most generous constant for the old bound, the new one overtakes it only once $\log_{10}(L/y)$ exceeds 334677.0. As an exponent statement the hierarchy is a real advance; as a numerical bound it improves nothing a finite orbit can reach.

The dense source floor at the restored input. The round derives the exponent the previous rounds had published without a formula: $y_1 \gtrsim N^{(\theta-\alpha)/(1-\alpha)}$. At ρ_c it is positive at all ten levels — the six $\hat{\alpha}_m$ and the four inherited exponents 1/6, 1/9, 4/45 and 1373/25856 — because θ_c exceeds every one, the binding case $\alpha = 1/6$ needing $\rho < 5$; and it falls at every level, as a lower bound must.

5. The dense branch is empty

The thirteenth round proves something different in kind: no fixed $\delta > 0$ and unbounded sequence of N have a completed B-support count $M_N \geq \delta N$. The chain is finite at every link — a mod-12 source floor, two duration floors, a pigeonhole localization, convexity on the origin gaps, a source corridor — and each was checked on its own (`COMPUTED_HERE`). Its one arithmetic input is the local best-approximation bound $\|q\beta\| > 1/((M_\beta(N)+2)q)$, decided from the exact continued fraction of β for every q below three scales, 111000 values: no violation. Its duration floors use the Diophantine input, but its conclusion needs it only through $\kappa_{13} < 1$, which is the statement $\theta \neq 0$ and so holds for any finite exponent; at ρ_c, $\kappa_{13} = 26183689/27183689$.

The round's object does not occur on a real orbit: among the first-crossing intervals of the orbits scanned there is no B-injection at all. The next round explains why.

6. Sparse support, and the price of anything denser

The fourteenth round splits the B-start count into completed and active intervals and closes both escapes it can see; if the envelope, slack and continued-fraction scales are all subpolynomial, the whole renewal process satisfies $R_N \leq N^{4/5+o(1)}$, and anything denser must pay one of three named prices: an A-envelope that rises, partial quotients $M_\beta(N) \geq N^{\chi(\kappa)}$, or slack $\delta_N \geq N^{\zeta(\kappa)}$, with $\chi(\kappa) = (5\kappa - 4)/3$ and $\zeta(\kappa) = (\kappa+1)/3$. The trichotomy rests on the identity $2\kappa - 1 - \chi(\kappa) = \zeta(\kappa)$, checked over random rational parameters rather than at the round's own value. The threshold 4/5 does not depend on ρ; the active-backlog exponent $1-\theta$ does, and is 4117/5117 at ρ_c.

Every suffix minimum steps up. The round's Theorem 3.1 — a suffix minimum $Y_s > 1$ of an injective orbit has $q_{s+1} = 1$, hence $Y_s \equiv 7$ or $11 \pmod{12}$ for $s \geq 1$ or a start prime to 3 — is about every suffix minimum, so real orbits test it: 16251 suffix minima across 2295 orbit windows, no violation (`COMPUTED_HERE`). And every one of them is an A-renewal. That is forced: a true suffix minimum with a first crossing is by definition a B-injection, so finding none of the one is finding none of the other — the reason the previous round's population was empty.

7. Records pinned by slack

The fifteenth round bounds the record count R_N from both sides by slack, $2^{-\Delta_N} N^{1-o(1)} \leq R_N \leq 2^{\delta_N} N^{o(1)}$, with Δ_N the largest slack so far, and concludes $\Delta_N + \delta_N \geq (1-o(1))\log_2 N$ — a Collatz-specific replacement for record theory, which alone allows records as sparse as $\sqrt{\log_2 N}$. None of this depends on ρ.

Its §10 needs no premise at all: on any segment of g steps from s, the number of valuations equal to 1 satisfies $N_1(s,g) \geq (2-\beta)g + (\delta_{s+g} - \delta_s)$, because every other valuation is at least 2. Checked on 127813 segments rooted at suffix minima and, as a control that the premise is truly unnecessary, on 59130 rooted anywhere: no violation, and the tightest case attained with slack exactly 0, so the bound

cannot be passed by accident (COMPUTED_HERE). The round's descent theorem, by contrast, bounds a total downward variation of slack at records that a convergent orbit does not have — 8447 record edges, all ascending — and the round says so itself.

8. The correction bank is an integer identity

The sixteenth round rests its central structural claim on the programme's correction bank $\mathcal{A}_n = 2^{-\delta_n} Y_n$ and its increment $\mathcal{A}_{n+1} - \mathcal{A}_n = \frac{1}{3} 2^{-\delta_n}$. Written so, checking it needs β. It does not: since $2^{\beta n} = 3^n$, the bank is $2^{K_n} Y_n / 3^n$, and the increment is the identity $2^{K_{n+1}} Y_{n+1} - 3 \cdot 2^{K_n} Y_n = 2^{K_n}$, immediate from $Y_{n+1} = (3Y_n + 1)/2^q$ — checked on 123005 steps, no violation (COMPUTED_HERE). Consecutive suffix-minimum gaps, unlike the objects of the previous rounds, are everywhere on finite windows: 4069 gaps of length at least 2, on which every claim of the round holds — the next record below every interior state, the ratio cap, interior slack domination, and a fully suffix-supercritical tail on 11305 suffixes. Its landing phases are residue arithmetic and hold on 3057 endpoints of one phase and 1012 of the other, and both explicit bridges it publishes, from 71 and from 223, rebuild exactly from the map. The forced gap exponent 1/5 does not depend on ρ.

9. What real orbits can and cannot test here

The rounds' theorems fall into three kinds, and this paper keeps them apart. **Identities about any orbit** — laminarity, the annulus identity, the correction bank, the valuation-count bound — were tested on real orbits and hold, several of them attained. **Statements whose objects real orbits supply** — suffix minima, consecutive gaps — were tested on populations in the thousands. **Statements about the hypothetical divergent branch** — B-injections, record descents, long segments above their source — have populations that are empty or nearly so on real orbits, and were checked as algebra and as implications between finite quantities; a zero violation count over an empty population is reported as the empty population it is.

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

Each gate was drilled on its own, defects planted one at a time and each required to turn red the check named for it. 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. The gate that re-evaluates the exponents found one fault of its own on its first drill: with the review's exponent unreadable, it fell back to a default of zero, divided by it, and died without a log — red, but a crash, and a crash is not a catch. It now refuses with a log that says why.

11. 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 programme's rounds this paper checks: the survival closure, the congestion rigidity, the transport hierarchy, the depth collision, the sparse-support exhaustion, the record sparsity and the record-gap transport.

Neo.K set the programme's target — a divergent orbit read through its slack and its renewals — and wrote its rounds with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — rechecked every round in exact arithmetic, tested the structural core on real orbits, found the premise stated weaker than used and the constants that outgrow the exponent, read the review behind the Diophantine input and restored it at the value the review supports, re-evaluated every exponent derived from it, built and drilled the gates, and wrote this paper. The rounds' mathematics is the programme's; nothing in it 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 exact verification of the rounds, the separation of what real orbits test from what they cannot, and the reading of the programme's one external input from its source.

12. What is open

1. The Collatz conjecture, and whether a divergent orbit's renewals can be sparser than the rounds allow.
2. The exact exponent of the Wu–Wang linear form, and whether its "sufficiently large" is effective (`UNMEASURED`).
3. The asymptotic halves of the rounds' statements, and the constant C_ε they consume.

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 survey is Fan, Queffélec and Queffélec's; the published review is zbMATH Open's, of Wu and Wang's paper.

1. Aihua Fan, Hervé Queffélec and Martine Queffélec, *Old and new results on the Furstenberg sets*, arXiv:2303.06850 (2023).
2. Georges Rhin, *Approximants de Padé et mesures effectives d'irrationalité*, in Séminaire de Théorie des Nombres, Paris 1985–86, Progress in Mathematics, Birkhäuser, 1987, 155–164, doi:10.1007/978-1-4757-4267-1_11.
3. 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.
- The mpmath development team, *mpmath: a Python library for arbitrary-precision floating-point arithmetic* (version 1.3.0), https://mpmath.org/ (2023).

A. Reproduction

Everything is Python standard library, no network, and every inequality linear in β is decided by comparing integers. Each gate writes a JSON log; the review of the paper behind the Diophantine input is archived byte for byte with its digest. 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 verification ledger**, rendered from the line's results file: https://verify.evemisslab.com/collatz-verification-zhuiheng/;
- **the round-by-round record**: https://amral.evemisslab.com/collatz/recompute/.

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`).