Polylogarithmic Descent for Almost All Collatz Orbits in Natural Density

Idris Ali Shaik

Independent researcher, Dubai, United Arab Emirates

Email: shaikidris@apache.org

ORCID: 0009-0009-9699-9712

Version: 3.2.3 - preprint

Date: 17 August 2026

Preprint DOI: 10.5281/zenodo.21984038

Lean 4 software DOI: 10.5281/zenodo.21930432

2020 Mathematics Subject Classification: Primary 11B83; Secondary 37P99, 60G40

Keywords: Collatz map, almost all Collatz orbits, natural density, first passage, stopping map, parity vectors, polylogarithmic descent

Abstract

We prove that, in ordinary natural density, almost every positive integer n admits a shortcut Collatz iterate of polylogarithmic size within $O(\log n)$ steps, with every iterate up to that time at most $n^{1+\beta}$ for any fixed $\beta > 0$. The argument supplies an explicit critical polylogarithmic exponent, quantitative exceptional-set rates, and a logarithmic witnessing clock.

Let T be the shortcut Collatz map, with $T(n) = n/2$ for even n and $T(n) = (3n+1)/2$ for odd n. Put $\varkappa_* = 1 - H_2(\log_3 2)$, $A_{\mathrm{FP}} = (2\varkappa_*)^{-1} = 9.9911133419\ldots$, and $c_* = 2/\log(4/3) = 6.9521189935\ldots$, where H_2 is binary entropy. For every fixed $A > A_{\mathrm{FP}}$, $c > c_*$, $\beta > 0$, and $0 < \gamma < \varkappa_*(A - A_{\mathrm{FP}})$, there is $C_{\mathrm{tar}} > 0$ such that, as $X \to \infty$, all but $O_{A,c,\beta,\gamma}(X/(\log X)^{\gamma})$ integers $n \le X$ admit an integer $k < c\log n$ satisfying

$$T^k(n) \le C_{\mathrm{tar}}(\log n)^A, \qquad \max_{0\le j\le k} T^j(n) \le n^{1+\beta}.$$

At the critical principal exponent A_{FP}, the same clock and ceiling hold after adding explicit log-log factors: for every fixed $D > 0$, with a possibly different C_{tar}, the target is

$$C_{\mathrm{tar}}(\log n)^{A_{\mathrm{FP}}}(\log\log n)^{2A_{\mathrm{FP}}}(\log\log\log n)^{D},$$

with exceptional count $O_{D,c,\beta,\gamma}(X/(\log\log\log X)^{\gamma})$ whenever $0 < \gamma < D\varkappa_*$. Neither the pure target $C(\log n)^{A_{\mathrm{FP}}}$ nor a bounded final multiplier at the critical secondary scale is asserted. For every fixed $0 < \delta < 1$, $c > c_*$, and $\beta > 0$, the weaker target $\exp((\log n)^{1-\delta})$ is reached with exceptional count $O_{\delta,c,\beta}(X\exp(-\gamma_{\delta,c,\beta}(\log X)^{1-\delta}))$ for some $\gamma_{\delta,c,\beta} > 0$, with the same shortcut clock and orbit ceiling. For the un-accelerated Collatz map, whose odd step is $n \mapsto 3n+1$, every clock constant greater than $3/\log(4/3)$ is allowed.

On each dyadic shell $[2^M, 2^{M+1})$, exact parity-word counting controls prefixes; a terminal odd-step tail controls small blocks that miss their next threshold in time. Decreasing thresholds make later failures direct first passages of the original orbit, so only $O(\sqrt{M\log M})$ cumulative times are counted instead of a linear horizon.

1. Introduction and main results

We ask how small a value can be forced along the Collatz orbit of almost every initial integer in ordinary natural density, while retaining a logarithmic witnessing time and a bound on the orbit before that witness.

For classical background and broad surveys of the $3x+1$ problem, see Crandall [1], Lagarias [7], and Wirsching [13].

We work throughout with the shortcut Collatz map

$$T(n) = \begin{cases} n/2, & n \equiv 0 \pmod 2, \\ (3n+1)/2, & n \equiv 1 \pmod 2. \end{cases}$$

One application of T is one shortcut step, and we write $T_{\min}(n) = \min_{k \geq 0} T^k(n)$. All unqualified logarithms are natural, $\log_2$ denotes the base-two logarithm, and $\mathbb{N} = \{1, 2, 3, \dots\}$. For $M \geq 0$, let

$$I_M = [2^M, 2^{M+1}) \cap \mathbb{N}$$

be the M-th dyadic shell; we call M its shell rank. A set $S \subseteq \mathbb{N}$ has natural density one if

$$\frac{\#\{1 \leq n \leq X : n \notin S\}}{X} \longrightarrow 0,$$

or, equivalently, if $\#(S \cap [1, X])/X \to 1$. Put

$$a_0 = \frac{\log_2 3}{2}. \tag{1.1}$$

Define binary entropy by

$$H_2(p) = -p \log_2 p - (1-p) \log_2 (1-p),$$

and put

$$p_* = \log_3 2, \qquad A_{FP} = \frac{1}{2(1 - H_2(p_*))}, \qquad c_* = \frac{2}{\log(4/3)}. \tag{1.2}$$

The subscript FP records that this exponent comes from the first-passage argument.

Where the constants come from

The two constants in (1.2) record the two quantitative costs of the proof. First,

$$1 - a_0 = \frac{\log_2(4/3)}{2}, \qquad c_* = \frac{1}{(1-a_0)\log 2}.$$

Thus $1 - a_0$ is the mean base-two logarithmic loss of the multiplicative main term per shortcut step, and $c_* \log n$ is the corresponding mean-drift clock for clearing $\log_2 n$ bits. This interpretation does not assert a matching lower bound or optimality theorem. Drift heuristics and related stochastic models are classical; see Kontorovich–Lagarias [4] and Lagarias–Weiss [8]. The estimates below nevertheless come from exact parity-word counting, not from a stochastic model of an individual orbit. Second, put

$$\varkappa_* = 1 - H_2(p_*).$$

The terminal odd-step count on a shell has an exact binomial law. Suppose the iteration stops at a shell rank $L = L(M)$. The upper tail at the critical proportion p_* has exponential rate $\varkappa_* \log 2$ and prefactor of order $L^{-1/2}$. The proof also sums over $O(\sqrt{M \log M})$ possible first-passage times and pays a factor of order L for the reverse loss. Together these terms leave the exponent margin

$$\varkappa_* L - \frac{1}{2}\log_2 M - \log_2 \log M .$$

Its divergence is the endpoint condition in the main theorem. For $L \sim A \log_2 M$, the leading power changes sign at $A\varkappa_* = 1/2$, giving $A_{FP} = 1/(2\varkappa_*)$. At equality the remaining logarithmic margin forces an additional slowly growing factor. Thus A_{FP} is a threshold of the present argument rather than an intrinsic constant of the Collatz map.

The main theorem permits a bounded exponent sequence (A_M), fixed in advance, with one exponent for each shell I_M. This single formulation includes every fixed exponent $A > A_{FP}$ and also exponents that approach A_{FP}. The exact admissibility condition is part of the statement; its most useful explicit specializations follow immediately afterward.

Theorem 1.1 (shell-dependent polylogarithmic descent)

Let $(A_M)_{M \ge 0}$ be a bounded real sequence and put

$$L_M = \lceil A_M \log_2(M+2) \rceil, \qquad \Delta_M = \varkappa_* L_M - \frac{1}{2}\log_2(M+2) - \log_2 \log(M+3). \tag{1.3}$$

Assume

$$\Delta_M \longrightarrow +\infty . \tag{1.4}$$

For every fixed $c > c_*$ and $\beta > 0$, there are constants $C_{tar}, C_{exc}, \varepsilon > 0$ and an integer $M_0 \ge 1$ such that the following holds. These constants may depend on the bounded sequence (A_M), as well as on c and β, but are independent of M. Let E_M be the set of $n \in I_M$ for which no integer $0 \le k < c \log n$ simultaneously satisfies

$$T^k(n) \le C_{tar}(\log n)^{A_M}, \qquad \max_{0 \le j \le k} T^j(n) \le n^{1+\beta}.$$

Then, for every $M \ge M_0$,

$$\boxed{\frac{\# E_M}{2^M} \le C_{exc}\left(2^{-\Delta_M} + M^{-\varepsilon}\right).} \tag{1.5}$$

Consequently these simultaneous witnesses exist on a set of natural density one.

The ceiling in L_M changes Δ_M by only $O(1)$. Hence, when $A_M \to A_{FP}$, condition (1.4) is equivalently

$$\varkappa_*(A_M - A_{FP}) \log M - \log\log M \longrightarrow +\infty . \tag{1.6}$$

Corollary 1.2 (explicit descent scales)

Fix $c > c_*$ and $\beta > 0$. Each numbered statement below has its own constant $C_{tar} > 0$, allowed to depend on all parameters displayed in that statement. The landing occurs before $c \log n$ shortcut steps, and its witness satisfies $\max_{j \le k} T^j(n) \le n^{1+\beta}$.

1. For every fixed $A > A_{FP}$ and every $0 < \gamma < \varkappa_*(A - A_{FP})$, the target

$$C_{tar}(\log n)^{A}$$

fails for at most $O_{A,c,\beta,\gamma}(X/(\log X)^{\gamma})$ integers $n \le X$.

2. For every fixed $B > 2A_{FP}$ (equivalently $B\varkappa_* > 1$) and every $0 < \gamma < B\varkappa_* - 1$, the target

$$C_{tar}(\log n)^{A_{FP}}(\log\log n)^{B}$$

fails for at most $O_{B,c,\beta,\gamma}(X/(\log\log X)^{\gamma})$ integers $n \le X$.

3. For every fixed $D > 0$ and every $0 < \gamma < D\varkappa_*$, the target

$$C_{tar}(\log n)^{A_{FP}}(\log\log n)^{2A_{FP}}(\log\log\log n)^{D} \tag{1.7}$$

fails for at most $O_{D,c,\beta,\gamma}(X/(\log\log\log X)^{\gamma})$ integers $n \le X$.

More generally, for any prescribed function $\Omega : [1,\infty) \to (0,\infty)$ with $\Omega(x) \to \infty$, there is an eventually nondecreasing function $\tilde{\Omega}$ such that

$$1 \le \tilde{\Omega}(x) \le \Omega(x) \quad \text{eventually}, \qquad \tilde{\Omega}(x) \to \infty, \qquad \tilde{\Omega}(x) = x^{o(1)},$$

and the final factor in (1.7) may be replaced, on a natural-density-one set, by $\tilde{\Omega}(\log\log n)$. Consequently the conclusion also holds with the larger factor $\Omega(\log\log n)$. No monotonicity is required of Ω. The constants and the retained set may depend on Ω, c, β.

Every numbered target above, and the moving $\tilde{\Omega}$-target constructed above, is $o(n)$, so it is a genuine descent for all sufficiently large n. The assertions are density-one rather than pointwise. Neither the landing constant C_{tar} nor the onset of the asymptotic comparison is made effective.

Theorem 1.3 (quantitative stretched-logarithmic descent)

For every fixed $0 < \delta < 1$, $c > c_*$, and $\beta > 0$, there are constants $C_{\delta,c,\beta}, \gamma_{\delta,c,\beta}, X_{\delta,c,\beta} > 0$ with the following property. For every $X \ge X_{\delta,c,\beta}$, let $E_{\delta,c,\beta}(X)$ be the set of integers $1 \le n \le X$ for which no integer $0 \le k < c\log n$ simultaneously satisfies

$$T^{k}(n) \le \exp((\log n)^{1-\delta}), \qquad \max_{0\le j\le k} T^{j}(n) \le n^{1+\beta}.$$

Then

$$\# E_{\delta,c,\beta}(X) \le C_{\delta,c,\beta}X\exp(-\gamma_{\delta,c,\beta}(\log X)^{1-\delta}). \tag{1.8}$$

The endpoint $\delta = 1$ is not asserted.

Corollary 1.4 (raw clock)

Put

$$\mathrm{Col}(n) = \begin{cases} n/2, & n \equiv 0 \pmod 2, \\ 3n+1, & n \equiv 1 \pmod 2. \end{cases}$$

For every fixed $\beta > 0$ and

$$c_{raw} > \frac{3}{\log(4/3)}, \tag{1.9}$$

the conclusions and exceptional-set estimates in Theorem 1.1, Corollary 1.2, and Theorem 1.3 remain valid for Col, with an integer $\ell < c_{\mathrm{raw}} \log n$ and the raw-orbit ceiling

$$\max_{0 \le j \le \ell} \mathrm{Col}^{j}(n) \le n^{1+\beta}.$$

All parameter ranges remain exactly as stated in those results; their existential constants may additionally depend on c_{raw}. The same raw index ℓ supplies both the corresponding landing and this ceiling. Since $3/\log(4/3) = 10.42817849\ldots$, the explicit constant 10.44 remains admissible.

Corollary 1.5 (fixed powers and the independent graded-clock companion)

For every fixed $\alpha > 0$, the exceptional set for $T_{\min}(n) > n^{\alpha}$ is

$$O_{\alpha,\sigma}(X \exp(-c_{\alpha,\sigma}(\log X)^{\sigma})) \tag{1.10}$$

for every fixed $0 < \sigma < 1$, where $c_{\alpha,\sigma} > 0$.

If $0 < \alpha < 1$, then for every fixed $\varepsilon > 0$ the same fixed-power target is also reached on a natural-density-one set: for every sufficiently large retained n, there is an integer $k \ge 0$ with $T^{k}(n) \le n^{\alpha}$ and

$$k < \left(\frac{2(1-\alpha)}{\log(4/3)} + \varepsilon \right) \log n \tag{1.11}$$

shortcut steps. Thus the clock decreases continuously with the amount of logarithmic height that remains: for example, the limiting coefficient at $\alpha = 1/2$ is $1/\log(4/3) = 3.47605949\ldots$, rather than the full descent coefficient c_*.

Terminology and recurring notation

We use the following terms consistently. A *threshold* is a dyadic value 2^{q} that the orbit seeks to cross from above. Its *first passage* is the first iterate at or below that threshold, and the value reached there is the *landing*. Successive landings that start new passages are called *checkpoints*; together they form a decreasing *threshold chain*.

A *certificate* is a condition on the parity data that guarantees the stated orbit bounds. At high shell ranks the certificate controls every parity prefix and hence every intermediate iterate. At low ranks the proof instead uses a *timeout*: failure to cross the next threshold within the allotted number of steps. The *switch rank* separates these high and low phases. *Transport* means counting sources in the original shell whose first-passage landing lies in a specified target set.

A *witnessing time* is an iterate at which the asserted descent has occurred. The *same-witness orbit ceiling* bounds every earlier iterate through that same time; it is not a statement about the remainder of the orbit. The global symbols $T, \mathrm{Col}, T_{\min}, I_M, a_0, p_*, \varkappa_*, A_{\mathrm{FP}}$, and c_* retain the meanings fixed above. Symbols used only inside one proof stage are introduced at the start of the section where they are needed; in particular, the high- and low-phase parameters are collected at the beginning of Section 6.

Proof architecture

The proof has two sharply separated parts and six load-bearing steps. Its only probabilistic input is exact counting of parity words on a dyadic shell; the first-crossing estimates and the assembly of the decreasing threshold chain are deterministic.

Exact probabilistic certification.

1. Every parity word of length M occurs exactly once on the shell I_M.

2. Counting words with an atypical prefix gives an exponentially small exceptional set and bounds the orbit at every time in the block.

Deterministic Collatz transport and assembly.

3. First-passage reversal bounds the number of sources for each fixed passage time and landing value, while retaining the accumulated effect of the additive $+1$ terms.
4. Along a decreasing threshold chain, every later landing is also a direct first passage from the original source shell.
5. Rank-dependent high barriers confine every time at which certification can first fail to a common set of size $O(\sqrt{M \log M})$.
6. Summing only over the feasible passage times and using the critical tail for small-shell blocks that do not cross in time give the exact margin Δ_M. Its positive divergence yields the principal exponent A_{FP}.

The time-support bound is stated and proved in Lemma 6.1. Proposition 6.4 assembles the critical shell estimate, and (6.20)–(6.21) convert it to the margin Δ_M explained in the opening constant discussion.

Steps 3 and 4 are the shared deterministic part of the proof. Applying them with one fixed barrier gives the stretched-logarithmic exceptional rate. Applying them with shrinking large-scale bounds and the small-scale timeout test gives both the fixed polylogarithmic targets and the shell-dependent exponent in the headline.

The diagram below shows this division and the flow from one dyadic shell to successively lower thresholds.

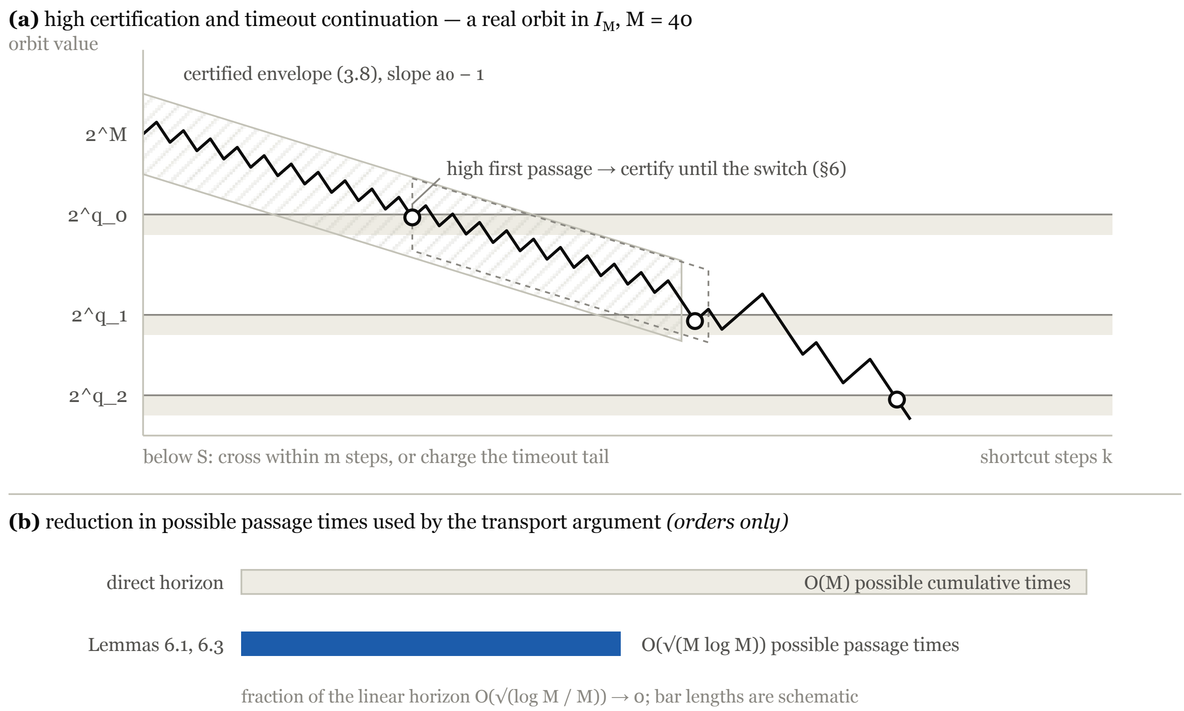

The proof architecture. Panel (a) follows an actual shortcut-Collatz orbit inside the certified high-rank envelope and marks successive threshold landings; below the switch, the timeout rule replaces low-rank certification. Panel (b) records the proved reduction in the number of possible cumulative passage times; bar lengths are schematic and do not encode the unspecified constant in Lemmas 6.1 and 6.3.

Relation to previous almost-all results

Terras [11] and Everett [2] proved the classical first-descent statement for almost every starting value; Terras's follow-up [12] addressed the associated density question.

The earlier endpoint-transport preprint [9] proves stretched-logarithmic descent in the range $0 < \delta < 0.251245530155874\ldots$ by fixed-total endpoint fibers, a Rényi estimate, and endpoint-only iteration. The present proof is logically independent of that theorem. Its proof architecture and principal conclusions are substantially different: high-rank all-prefix certification, first-passage reversal along decreasing thresholds, and a small-rank timeout tail yield fixed polylogarithmic targets above the explicit critical exponent, a moving critical endpoint, and a same-witness orbit ceiling. Theorem 1.3 also gives an independent stretched-logarithmic companion for every $0 < \delta < 1$, extending the parameter range obtained in [9].

The main obstruction is not the contraction of one typical block, but the transport of sparse certification failures through a generated sequence of landings. Direct aggregation over every time up to the horizon incurs a linear loss. The proof instead organizes long multi-landing passages in natural density without that linear time-union loss. Step 5 above reduces the possible times to $O(\sqrt{M \log M})$ possibilities; this linear-to-square-root reduction is the quantitative pivot of the paper.

The published results of Korec and Tao, the preprint of Inselmann, and two recent Tao-to-natural-density bridge manuscripts (J. Allikvere, unpublished observations; L. Mazur, unpublished observations) are compared below only on the coordinates used in their stated theorems. Publication status and proof lineage are recorded in the surrounding prose rather than in the table.

Work	Density	Target reached for almost all starts	Clock or quantitative feature
Korec [5]	natural	n^{θ}, every fixed $\theta > a_0$	no clock or quantitative exception used here
Inselmann [3]	natural	n^{ε}, every fixed $\varepsilon > 0$	$2 \log n / \log(4/3)$ shortcut steps; density convergence
Tao [10]	logarithmic	every $f(n) \to \infty$	no single global clock in the headline theorem; logarithmic-density estimate
Mazur	natural, bridged from Tao's logarithmic-density framework	every $f(n) \to \infty$	$< 436 \log n$ unaccelerated steps; fixed-target logarithmic rate
Allikvere	natural, bridged from Tao's logarithmic-density framework	every $f(n) \to \infty$	$< 12 \log n$ unaccelerated steps; $O((\log N_0)^{-1/29} + X^{-1/2000})$ for a fixed target
This paper	natural	$C_{tar}(\log n)^{A}$ for every fixed $A > A_{FP}$; moving critical endpoint in Corollary 1.2	every shortcut constant $> c_*$; every unaccelerated constant $> 3/\log(4/3)$; target-dependent quantitative rates and an orbit-height bound through the witnessing time

Theorem 1.1 additionally bounds every preceding iterate through the same witnessing time: $\max_{j \le k} T^{j}(n) \le n^{1+\beta}$. This bound is not part of the headline statements in those unpublished observations; this is a comparison of stated outputs, not a claim that their proof-internal trajectory estimates cannot yield related bounds.

The moving target of Theorem 1.1 includes the critical specialization $C_{tar}(\log n)^{A_{FP}}(\log\log n)^{2A_{FP}}\tilde{\Omega}(\log\log n)$, where $\tilde{\Omega}(x) = x^{o(1)} \to \infty$ may be chosen eventually below each prescribed $\Omega(x) \to \infty$. It is smaller than every fixed power and every fixed stretched-logarithmic target in Theorem 1.3, but it still contains the fixed divergent core $(\log n)^{A_{FP}}(\log\log n)^{2A_{FP}}$. It is therefore weaker than an arbitrary diverging function. At the target-only level, Tao's theorem specializes to every target displayed here, albeit in logarithmic rather than natural density. Thus the present theorem does not supersede the arbitrary-threshold conclusions reported in those unpublished observations. Conversely, those target statements do not supply the present natural-density rates, logarithmic clock, or orbit-height bound, and the clock scale already appears in [3]. These axes are deliberately not collapsed into a single ordering.

Complementary inverse-tree work obtains lower bounds for the number of integers whose orbits reach 1 by means of difference inequalities; see Krasikov–Lagarias [6]. That direction is distinct from the forward first-passage transport used here.

Allikvere conditions Tao's Syracuse/Fourier inputs on the total valuation and supplies a uniform-measure transfer. Mazur starts from the same logarithmic-density frontier and adds a quantitative phase-gap and two-adic lift. Their shared bridge architecture differs from the exact Boolean-cube certification and deterministic first-passage transport separation displayed above.

2. Density, parity, and affine iterates

The argument uses three common tools: a shell-to-prefix summation rule, exact parity coding on each dyadic shell, and an affine formula for the iterate. This section establishes all three before they are used.

For $S \subseteq \mathbb{N}$, put

$$B_S(X) = \#\{1 \le n \le X : n \notin S\}.$$

We call S (C,D)**-dense** if

$$B_S(X) \le CX^{1-D} \qquad (X \ge 1), \tag{2.1}$$

where $C > 0$ and $0 < D \le 1$. Every such set has natural density one.

We use the following varying-rate shell summation.

Lemma 2.1 (varying-rate dyadic summation)

Let $0 < \sigma \le 1$, $A, \gamma > 0$, and suppose that, for all sufficiently large M,

$$\#(E \cap [2^M, 2^{M+1})) \le Ae^{-\gamma(M+4)^\sigma}2^M. \tag{2.2}$$

Then there are $\gamma' > 0$ and X_0 such that

$$\#(E \cap [1,X]) \le (2A+1)Xe^{-\gamma'(\log X)^\sigma} \qquad (X \ge X_0). \tag{2.3}$$

More generally, if a nonnegative sequence $e_M \to 0$ satisfies

$$\#(E \cap I_M) \le e_M 2^M \tag{2.3a}$$

for all sufficiently large M, then E has natural density zero.

Proof

Let $2^J \le X < 2^{J+1}$. The shells below $J/2$ contain fewer than $X^{1/2}$ integers. On the remaining shells,

$$M + 4 \ge \frac{\log X}{4\log 2}.$$

Thus their total contribution is at most

$$2AX\exp\left(-\frac{\gamma}{(4\log 2)^{\sigma}}(\log X)^{\sigma}\right).$$

Choose

$$\gamma' = \min\left\{\frac{\gamma}{(4\log 2)^{\sigma}}, \frac{1}{2}\right\}.$$

For all sufficiently large X, one has $X^{1/2} \le Xe^{-\gamma'(\log X)^{\sigma}}$. Adding the early and late contributions proves (2.3).

For the final assertion, fix $\epsilon > 0$ and choose M_0 so that $e_M \le \epsilon$ for $M \ge M_0$. The finitely many earlier shells contribute $o(X)$. If $2^J \le X < 2^{J+1}$, the remaining shells contribute at most

$$\epsilon \sum_{M=M_0}^{J} 2^M \le 2\epsilon X.$$

Letting $\epsilon \downarrow 0$ proves natural density zero. □

For $n \ge 1$, define

$$p_i(n) = 1_{\{T^i(n)\ \mathrm{odd}\}}, \qquad s_k(n) = \sum_{i=0}^{k-1} p_i(n). \tag{2.4}$$

The exact parity coding below is classical in the stopping-time literature; see, for example, Terras [11]. Its short proof is included because no external result is needed by the argument.

Proposition 2.2 (parity-vector bijection)

For every $M \ge 0$,

$$n \bmod 2^M \mapsto (p_0(n), \ldots, p_{M-1}(n)) \tag{2.5}$$

is a bijection from $\mathbb{Z}/2^M\mathbb{Z}$ to $\{0,1\}^M$.

Proof

The relation

$$n \equiv n' \pmod{2^{k+1}} \quad \Longrightarrow \quad T(n) \equiv T(n') \pmod{2^k}$$

shows inductively that the first M parity bits depend only on $n \bmod 2^M$. For the converse, induct on the word length. Suppose the first bit is e, and use the induction hypothesis to recover the residue $m = T(n) \bmod 2^k$ from the remaining k bits. If $e = 0$, the unique compatible residue is

$$n \equiv 2m \pmod{2^{k+1}}.$$

If $e = 1$, it is

$$n \equiv (2m - 1)3^{-1} \pmod{2^{k+1}},$$

where 3 is a unit modulo 2^{k+1}. The first residue is even, the second is odd, and in either case applying T recovers $m \bmod 2^k$. Thus every word has exactly one preimage residue. □

Proposition 2.3 (exact affine iterate)

For every $n \ge 1$ and $k \ge 0$,

$$T^k(n) = \frac{3^{s_k(n)}}{2^k}n + \sum_{i=0}^{k-1} \frac{p_i(n)3^{s_k(n)-s_{i+1}(n)}}{2^{k-i}}. \tag{2.6}$$

Proof

The assertion is trivial at $k = 0$. Passing from k to $k + 1$, an even letter divides every existing term by two. An odd letter multiplies every existing term by $3/2$ and adds $1/2$. This is exactly the displayed update. □

3. A dense set with controlled parity prefixes

The parity bijection identifies a complete dyadic shell with the Boolean cube. We now discard the sparse words whose prefixes deviate too far from their mean and obtain the pathwise orbit envelope needed to guarantee the first passages used later.

For a parity word of length M, put

$$Y_k = 2s_k - k, \qquad H_M = \max_{0\le k\le M}\left|s_k - \frac{k}{2}\right| = \frac{1}{2}\max_{0\le k\le M}|Y_k|. \tag{3.1}$$

For $0 < p, q < 1$, write

$$D(p\|q) = p\log\frac{p}{q} + (1-p)\log\frac{1-p}{1-q}$$

for binary relative entropy. For $0 \le t < 1/2$, put

$$I(t) = D\left(\frac{1}{2}+t \,\middle\|\, \frac{1}{2}\right) = \left(\frac{1}{2}+t\right)\log(1+2t) + \left(\frac{1}{2}-t\right)\log(1-2t). \tag{3.2}$$

Lemma 3.1 (maximal deviation of parity prefixes)

If $M \ge 1$ and $0 \le h < M/2$, then

$$2^{-M}\#\{w \in \{0,1\}^M : H_M(w) > h\} \le 2\exp(-MI(h/M)). \tag{3.3}$$

In particular, the right side is at most $2\exp(-2h^2/M)$.

Moreover, if $0 < t_0 \le t \le t_1 < 1/2$, then there is a constant C_{t_0,t_1} such that, for every $M \ge 1$,

$$2^{-M}\#\{w : H_M(w) > tM\} \le C_{t_0,t_1}M^{-1/2}e^{-MI(t)}. \tag{3.3a}$$

Proof

Fix $\theta > 0$ and stop the Boolean tree at the first prefix u with $|Y(u)| > 2h$. For a frontier node of depth k, set

$$\Phi_\theta(u) = 2^{-k}\cosh(\theta Y(u))(\cosh\theta)^{M-k}.$$

The identity

$$\cosh(\theta(y-1)) + \cosh(\theta(y+1)) = 2\cosh(\theta y)\cosh\theta$$

shows that replacing an unexpanded node by its two children preserves the total potential. The frontier formed by first-crossing prefixes and noncrossing depth-M words therefore has total $(\cosh\theta)^M$. Each crossing prefix contributes at least $2^{-|u|}\cosh(2\theta h)$, whence

$$\Pr(H_M > h) \le 2\exp(M\log\cosh\theta - 2\theta h). \tag{3.4}$$

Take $u = 2h/M$ and $\theta = \operatorname{arctanh} u$. The elementary Legendre identity

$$u\theta - \log\cosh\theta = D\left(\frac{1+u}{2}\,\Big\|\,\frac{1}{2}\right) = I(h/M)$$

proves (3.3). Pinsker's inequality in this binary case, or the direct convexity estimate $I(t) \geq 2t^2$, gives the quadratic assertion.

It remains to prove (3.3a). Put $a = \lfloor 2tM \rfloor + 1$. Reflection, followed by symmetry between the two boundary signs, gives

$$\Pr(H_M > tM) \leq 4\Pr(Y_M \geq a - 1). \tag{3.3b}$$

This form is insensitive to the parity of a: replacing the terminal threshold by the next reachable lattice point only decreases the event. If

$$k_0 = \left\lceil \frac{M + a - 1}{2} \right\rceil, \qquad p_M = \frac{k_0}{M},$$

then the probability on the right of (3.3b) is $2^{-M} \sum_{k=k_0}^{M} \binom{M}{k}$. For $k \geq k_0$, consecutive terms have ratio

$$\frac{\binom{M}{k+1}}{\binom{M}{k}} = \frac{M-k}{k+1}.$$

Uniformly for $t_0 \leq t \leq t_1$, this ratio is at most a fixed number strictly below one once M exceeds a fixed startup. The tail is therefore at most a fixed multiple of its first term. Uniform Stirling inequalities, with p_M confined to a compact subset of $(0,1)$, give

$$2^{-M}\binom{M}{k_0} \leq CM^{-1/2}\exp\{-MD(p_M \,\|\, 1/2)\}.$$

Here $p_M = 1/2 + t + O(M^{-1})$, uniformly in the displayed compact interval. The derivative of $D(p \,\|\, 1/2)$ is bounded on a slightly larger compact interval, so

$$MD(p_M \,\|\, 1/2) = MI(t) + O(1).$$

Indeed, applying the two-sided Stirling bounds to $M!$, $k_0!$, and $(M - k_0)!$ produces the displayed entropy exponent and a prefactor $O((Mp_M(1 - p_M))^{-1/2})$; compactness makes this $O(M^{-1/2})$ uniformly. Absorbing this bounded error and the finite startup into C_{t_0,t_1} proves (3.3a). □

Set

$$\varrho := \frac{\sqrt{3}}{2} = 2^{a_0 - 1}.$$

Define

$$r_k(n) = \sum_{i=0}^{k-1} \frac{p_i(n)}{3^{s_{i+1}(n)} 2^{k-i}}, \qquad d_k = s_k - \frac{k}{2}. \tag{3.5}$$

By Proposition 2.3,

$$T^k(n) = \varrho^k n 3^{d_k} + (r_k(n) 3^{k/2}) 3^{d_k}. \tag{3.6}$$

Lemma 3.2 (uniform affine correction)

If $H_M \le h$, then for every $1 \le k \le M$,

$$r_k(n)3^{k/2} \le (2+\sqrt{3})(1-\varrho^k)3^h < (2+\sqrt{3})3^h. \tag{3.7}$$

Proof

For $0 \le i < k$, the barrier gives $s_{i+1} \ge (i+1)/2 - h$. Hence

$$r_k 3^{k/2} \le 3^h \sum_{i=0}^{k-1} \frac{3^{(k-i-1)/2}}{2^{k-i}} = \frac{3^h}{2} \sum_{j=0}^{k-1} \varrho^j,$$

which is (3.7). □

For $0 < \eta \le 1$, let W_η be the set of integers n such that

$$\varrho^k n^{1-\eta} \le T^k(n) \le \varrho^k n^{1+\eta} \qquad (0 \le k \le \lfloor \log_2 n \rfloor). \tag{3.8}$$

Define

$$b_{\mathrm{ent}}(\eta) = I\left(\frac{\eta}{\log_2 3}\right). \tag{3.9}$$

Proposition 3.3 (density of the controlled-prefix set)

Let $0 < \eta < 1 - a_0$. For every fixed $0 < c < b_{\mathrm{ent}}(\eta)$, there is $K_{\eta,c} > 0$ such that

$$\frac{\#(W_\eta^c \cap I_M)}{2^M} \le K_{\eta,c} e^{-cM} \qquad (M \ge 0). \tag{3.10}$$

Consequently W_η is $(K'_{\eta,c}, c/\log 2)$-dense after increasing the fixed constant.

Proof

Put $L_3 = \log_2 3$. Since $c < b_{\mathrm{ent}}(\eta)$, choose a fixed $0 < \lambda < 1$ with

$$c < I\left(\frac{\lambda\eta}{L_3}\right). \tag{3.11}$$

On I_M, take

$$h = \frac{\lambda\eta M}{L_3}. \tag{3.12}$$

Assume $H_M \le h$. Since $3^h = 2^{\lambda\eta M} \le n^\eta$, the multiplicative term in (3.6) is at least $\varrho^k n^{1-\eta}$ and at most

$$2^{-(1-\lambda)\eta M} \varrho^k n^{1+\eta}. \tag{3.13}$$

Lemma 3.2 bounds the additive term by

$$(2+\sqrt{3})3^{2h} = (2+\sqrt{3})2^{2\lambda\eta M}. \tag{3.14}$$

The unused part of the upper envelope is at least

$$\left(1 - 2^{-(1-\lambda)\eta M}\right) \varrho^M 2^{(1+\eta)M}. \tag{3.15}$$

Indeed, before this uniform replacement the slack is $(1 - 2^{-(1-\lambda)\eta M})\varrho^k n^{1+\eta}$. Over all $0 \le k \le M$ and $n \in I_M$, its minimum occurs at $k = M$ and at the lower shell endpoint $n = 2^M$, which gives (3.15). Its binary exponent exceeds that of (3.14), because

$$a_0 + \eta - 2\lambda\eta \ge a_0 - \eta > 0. \tag{3.16}$$

Thus (3.8) holds for every $k \le M$ once M is sufficiently large.

Proposition 2.2 and Lemma 3.1 now give

$$\frac{\#(W_\eta^c \cap I_M)}{2^M} \le 2\exp\left[-MI\left(\frac{\lambda\eta}{L_3}\right)\right] \le 2e^{-cM}$$

on all sufficiently large shells. Enlarge $K_{\eta,c}$ to absorb the finite startup. Summing the resulting geometric shell series proves the global density assertion. □

4. First-passage linear transport

The barrier identifies a sparse set of bad orbit values. To count the original sources that reach such values, we must transport an arbitrary set of bad landings back to the original shell. Exact reversal at the first crossing, together with a bound that retains the reverse-product loss, provides that bridge.

For real $Y > 1$, define

$$\tau_Y(n) = \min\{h \ge 0 : T^h(n) \le Y\} \tag{4.1}$$

when the set is nonempty.

Lemma 4.1 (first-passage band and reverse product)

Let $1 < Y < 2^M$, let $n \in I_M$, and suppose $\tau_Y(n) = h \ge 1$. Write $x_j = T^j(n)$, $y = x_h$, and let s be the number of odd values among $x_0, \dots, x_{h-1}$. Then

$$\frac{Y}{2} < y \le Y \tag{4.2}$$

and

$$n = \frac{2^h y}{3^s} \prod_{\substack{0 \le j < h \\ x_j \text{ odd}}} \left(1 - \frac{1}{2x_{j+1}}\right). \tag{4.3}$$

Consequently, whenever $h < 2Y$,

$$\left(1 - \frac{h}{2Y}\right)\frac{2^h y}{3^s} \le n \le \frac{2^h y}{3^s}. \tag{4.4}$$

Proof

The final crossing cannot be odd: otherwise $x_h = (3x_{h-1} + 1)/2 > x_{h-1} > Y$. Hence $x_{h-1} = 2y$, proving (4.2). Reversing an even step gives $x_j = 2x_{j+1}$, while reversing an odd step gives

$$x_j = \frac{2x_{j+1} - 1}{3} = \frac{2}{3}x_{j+1}\left(1 - \frac{1}{2x_{j+1}}\right).$$

Multiplication proves (4.3). Every odd factor occurs before the final crossing, so its following state is greater than Y. Therefore $\prod_i (1 - u_i) \geq 1 - \sum_i u_i$ gives a lower product bound $1 - s/(2Y) \geq 1 - h/(2Y)$; the upper bound is one. □

Write the reverse product in (4.3) as $\prod_{j<h}(1 - u_j)$, where

$$u_j(n) = \begin{cases} \dfrac{1}{2T^{j+1}(n)}, & T^j(n) \text{ odd}, \\ 0, & T^j(n) \text{ even}, \end{cases} \qquad E_Y(n) = Y \sum_{j=0}^{h-1} u_j(n). \tag{4.5}$$

The factor Y records the additive reverse loss at the target scale: E_Y/Y is an additive upper bound for the defect of the reverse product. It also makes concatenation compatible with changing thresholds. If a segment is measured locally at $Y' \geq Y$, then its contribution after rescaling to the final threshold Y is exactly Y/Y' times its local contribution to $E_{Y'}$. Thus segment losses add after rescaling, which is the coordinate used below to bound the total loss along a decreasing threshold chain. The elementary product inequality gives

$$1 - \frac{E_Y(n)}{Y} \leq \prod_{j=0}^{h-1} (1 - u_j(n)) \leq 1. \tag{4.6}$$

Lemma 4.2 (fixed-time landing fibers with bounded reverse loss)

Let $D \geq 0$, and suppose

$$\frac{D}{Y} \leq \frac{1}{3}. \tag{4.7}$$

For a fixed pair (h, y), the number of sources $n \in I_M$ satisfying

$$\tau_Y(n) = h, \qquad T^h(n) = y, \qquad E_Y(n) \leq D$$

is at most

$$1 + 3D\frac{2^M}{Y}. \tag{4.8}$$

Proof

Put $P(n) = \prod_{j<h}(1 - u_j(n))$. By Lemma 4.1, if two such sources had odd counts $s_1 < s_2$, put $A_i = 2^h y/3^{s_i}$. Then $A_1 \geq 3A_2$, while (4.3) and (4.6) give

$$n_1 = A_1 P(n_1) \geq 3A_2\left(1 - \frac{D}{Y}\right) \geq 2A_2 \geq 2n_2.$$

This is impossible in the half-open shell I_M. The argument uses only positive products and introduces no division by a reverse product. Fix the common odd count and put $A = 2^h y/3^s$. All sources lie in $[(1 - D/Y)A, A]$. Nonemptiness, (4.7), and $n < 2^{M+1}$ imply $A < (3/2)2^{M+1}$, so this interval has length less than

$$3D\frac{2^M}{Y}.$$

Any additional congruence or residue restriction can only reduce the number of integer sources in this interval; no full-lattice occupancy is assumed. This proves (4.8). □

Proposition 4.3 (transport of a landing set with bounded reverse loss)

Put $J_Y = (Y/2, Y] \cap \mathbb{N}$. Let $1 < Y < 2^M$. Under (4.7), every $B \subseteq J_Y$ satisfies the exact bound

$$\#\left\{ n \in I_M : \begin{matrix} \tau_Y(n) \le H, & T^{\tau_Y(n)}(n) \in B, \\ E_Y(n) \le D \end{matrix} \right\} \le H\left(1 + 3D\frac{2^M}{Y}\right)|B|. \tag{4.9}$$

Consequently,

$$\#\left\{ n \in I_M : \begin{matrix} \tau_Y(n) \le H, & T^{\tau_Y(n)}(n) \in B, \\ E_Y(n) \le D \end{matrix} \right\} \le H(1 + 3D)\frac{2^M}{Y}|B|. \tag{4.10}$$

The estimate remains valid after arbitrary source restriction.

Proof

Since $Y < 2^M$, every passage time is positive. Sum Lemma 4.2 over the at most $H|B|$ pairs to obtain (4.9). The contribution $H|B|$ is at most $H(2^M/Y)|B|$, which gives (4.10). □

5. Threshold chains, direct passage, and reverse loss

This section proves two deterministic facts used by both main arguments. First, sequential crossings of decreasing thresholds collapse to one direct first passage from the original shell. Second, their reverse losses telescope. Theorem 1.3 applies these facts through the fixed-barrier route; the next section retains the same two facts and changes how failure is detected at small shell ranks.

For example, suppose an orbit first enters below 2^{30}, and from that landing later first enters below 2^{27}. No earlier point of the original orbit was below 2^{30}, hence none was below the smaller threshold 2^{27}; the second landing is therefore the original orbit's first entry below 2^{27}. The first lemma iterates this elementary observation along the whole decreasing threshold chain.

Fix

$$a_0 < r < 1, \qquad 0 < \eta < r - a_0, \qquad 0 < c_\eta < b_{\mathrm{ent}}(\eta). \tag{5.1}$$

Increase a fixed startup rank $M_0 = M_0(r, \eta, c_\eta)$ whenever needed below. If $x \in W_\eta \cap I_m$, put $q = \lfloor rm \rfloor$. For all $m \ge M_0$,

$$T^m(x) \le \varrho^m x^{1+\eta} < 2^{(a_0+\eta)m+1+\eta} \le 2^q. \tag{5.2}$$

Thus $\tau_{2^q}(x) \le m$, and (3.8) bounds the whole block by $x^{1+\eta}$.

Fix integers $M \ge L \ge M_0$, and start with $n_0 = n \in I_M$. Put

$$m_i = \lfloor \log_2 n_i \rfloor, \qquad t_0 = 0.$$

We call the points n_i the checkpoints of the chain. If $m_i < L$, stop successfully. If $m_i \ge L$ but $n_i \notin W_\eta$, stop with a certification failure. Otherwise define

$$q_i = \lfloor r m_i \rfloor, \quad h_i = \tau_{2^{q_i}}(n_i), \quad n_{i+1} = T^{h_i}(n_i), \quad t_{i+1} = t_i + h_i. \tag{5.3}$$

The first-passage band gives

$$2^{q_i - 1} < n_{i+1} \le 2^{q_i}, \qquad m_{i+1} \in \{q_i - 1, q_i\}. \tag{5.4}$$

Since $q_i \le m_i - 1$,

$$m_{i+1} \le q_i \le r m_i, \qquad q_{i+1} \le q_i - 1 \tag{5.5}$$

whenever the next block is executed. Consequently

$$t_i \le \sum_{j<i} m_j \le M \sum_{j<i} r^j < \frac{M}{1-r}. \tag{5.6}$$

Lemma 5.1 (direct passage along decreasing thresholds)

For every executed block i,

$$t_{i+1} = \tau_{2^{q_i}}(n_0), \qquad n_{i+1} = T^{\tau_{2^{q_i}}(n_0)}(n_0). \tag{5.7}$$

Proof

The assertion is the definition when $i = 0$. Inductively, every state before t_i exceeds the preceding threshold and hence the strictly smaller 2^{q_i}. At time t_i, one has $n_i \ge 2^{m_i} > 2^{q_i}$. The local block stays above 2^{q_i} until its endpoint. Thus $t_i + h_i$ is exactly the first crossing of the new threshold by the original orbit. □

The proof uses only that the thresholds decrease strictly and that every executed block is a genuine first passage to its threshold. We will use this generic form in Section 6 for the mixed high-rank/timeout chain; no certification condition on the low checkpoints is required.

For $q \ge 1$, put

$$J_q = (2^{q-1}, 2^q] \cap \mathbb{N}, \qquad B_q = W_\eta^c \cap J_q. \tag{5.8}$$

The interval J_q lies in I_{q-1} apart from its upper endpoint. Proposition 3.3 therefore gives

$$\frac{|B_q|}{2^q} \le C_{\eta,c_\eta}(e^{-c_\eta q} + 2^{-q}). \tag{5.9}$$

Lemma 5.2 (total reverse loss at a threshold)

If blocks $0, \dots, i$ are executed before the first certification failure, then the direct first passage in (5.7) satisfies

$$E_{2^{q_i}}(n) < \frac{q_i + 2}{r}. \tag{5.10}$$

Proof

In block j, every state following an odd step is strictly larger than 2^{q_j}; the final crossing is even. Therefore each reverse-loss increment is less than 2^{-q_j-1}. Loss incurred in block j is initially measured relative to its own threshold 2^{q_j}. Expressing it in units of the final threshold 2^{q_i} multiplies that loss by $2^{q_i-q_j}$. Since the block has at most m_j steps, summing these rescaled contributions gives

$$E_{2^{q_i}}(n) < \frac{1}{2} \sum_{j \le i} m_j 2^{q_i - q_j}.$$

Now $m_j < (q_j + 1)/r$, and the distinct integer thresholds q_j are all at least q_i. Thus, for each $d = q_j - q_i \ge 0$, there is at most one block at rank $q_i + d$. The worst case occupies every such rank, and the identities

$$\sum_{d \ge 0} 2^{-d} = 2, \qquad \sum_{d \ge 0} d 2^{-d} = 2$$

give

$$E_{2^{q_i}}(n) < \frac{1}{2r} \sum_{d \ge 0} (q_i + d + 1) 2^{-d} = \frac{q_i + 2}{r}.$$

□

More generally, the same calculation applies to any finite chain of genuine first-passage blocks for which the threshold ranks are distinct and decreasing, each block has duration at most its parent rank m_j, and $q_j + 1 > r_* m_j$ for one fixed $r_* > 0$. Its scaled reverse loss at the final threshold rank q_i is then less than $(q_i + 2)/r_*$. This is the form used for the timeout chain in Proposition 6.4.

Let $\mathrm{Fail}_{M,L}$ denote the sources whose chain reaches a certification failure before it reaches rank below L.

Theorem 5.3 (fixed-barrier failure bound)

Under (5.1), there are an integer $M_0 = M_0(r, \eta, c_\eta) \ge 1$ and a constant $C > 0$ such that, for all integers $M \ge L \ge M_0$,

$$\boxed{\frac{\#\,\mathrm{Fail}_{M,L}}{2^M} \le C\left[e^{-c_\eta M} + M(L+1)(e^{-c_\eta L} + 2^{-L})\right].} \tag{5.11}$$

Every source outside this set has an integer $k < M/(1-r)$ such that

$$T^k(n) < 2^L, \qquad \max_{0 \le j \le k} T^j(n) \le n^{1+\eta}. \tag{5.12}$$

Proof

An initial failure has proportion $O(e^{-c_\eta M})$. Otherwise let the first failed certification be a generated landing in B_q. Every earlier block and the block producing that landing are certified, so Lemma 5.2 applies; the bad landing itself is not assumed certified. Lemma 5.1 makes it a direct first passage from I_M, at time at most $M/(1-r) + 1$.

Increase M_0, if necessary, so that

$$\frac{q+2}{r 2^q} \le \frac{1}{3} \tag{5.13}$$

for every $q \ge M_0$. In Proposition 4.3, the available horizon satisfies $H \le M/(1-r) + 1 \ll_r M$, while with $D = (q+2)/r$ one has $1 + 3D \ll_r q + 1$. Combining (4.10) with (5.9) therefore bounds the proportion of sources whose first failed landing has rank q by

$$CM(q+1)(e^{-c_\eta q} + 2^{-q}). \tag{5.14}$$

The two elementary tail estimates

$$\sum_{q \ge L} (q+1) e^{-c_\eta q} \ll (L+1) e^{-c_\eta L}, \qquad \sum_{q \ge L} (q+1) 2^{-q} \ll (L+1) 2^{-L}$$

prove (5.11).

Outside the failure set, integer ranks strictly decrease until a state of rank below L is reached. Equations (5.2), (5.6), and the monotonic decrease of block endpoints prove (5.12). □

6. Shrinking prefix bounds and compressed passage times

Theorem 5.3 shows what the shared direct-passage and reverse-loss facts yield with one fixed certificate and a linear time horizon. For the stronger shell-dependent exponent, we retain those facts but shrink the high-rank barrier and replace low-rank certification by a timeout tail. The resulting set of possible cumulative times is much smaller.

The parameters used below fall into four groups. This table is a roadmap; the displayed equations following it give the precise definitions and the order in which the constants are chosen.

Symbols	Role	Definition or constraint
A_M, L_M, Δ_M	terminal schedule and its remaining exponent margin	fixed in Theorem 1.1; $L_M \asymp \log M$ and $\Delta_M \to \infty$ in the main application
r_{hi}	threshold ratio at high ranks	initially $a_0 < r_{hi} < 1$; strengthened for the final clock in (6.18)
τ	cap on the high-rank tolerance	initially $0 < \tau < r_{hi} - a_0$; also chosen below β and a_0 in (6.18)
D_{hi}	coefficient of the shrinking high-rank tolerance	chosen large enough in Proposition 6.4 to make high failures polynomially small
C_{sw}, S_M	switch coefficient and switch rank	$S_M = \lceil C_{sw} \log(M+2) \rceil$; C_{sw} is chosen after D_{hi} so that $D_{hi}/\sqrt{C_{sw}} \le \tau$ and $L_M < S_M$ eventually
$\eta_{M,m}$	high-rank tolerance at parent rank m	derived from τ, D_{hi}, M, m in (6.3)
K_0, r_L, q_L	bounded-gap low timeout rule	$K_0 > 0$ is fixed, $r_L = 1 - K_0/(2L)$, and $q_L(m) = \lfloor r_L m \rfloor$
r_*	uniform ratio used only in the reverse-loss estimate	a fixed rational with $0 < r_* < r_{hi}$ and, eventually, $r_* < r_{L_M}$

Thus only $r_{hi}, \tau, D_{hi}, C_{sw}, K_0$, and r_* are choices made inside the argument; the remaining quantities are determined by those choices and the prescribed terminal schedule.

High-rank stopped chain and compressed passage times. Choose fixed high-rank data

$$a_0 < r_{hi} < 1, \qquad 0 < \tau < r_{hi} - a_0, \tag{6.1}$$

and choose the high switch and tolerance below. For constants $D_{hi}, C_{sw} > 0$, put

$$S_M = \lceil C_{sw} \log(M+2) \rceil, \tag{6.2}$$

and, at a high parent rank $m \ge S_M$, certify with

$$\eta_{M,m} = \min\left\{ \tau, D_{hi} \sqrt{\frac{\log(M+2)}{m}} \right\}. \tag{6.3}$$

The high-rank stage construction at tolerance τ remains valid after lowering its tolerance to $\eta_{M,m}$; the density of the resulting rank-dependent certification sets will be estimated separately below.

Fix M, write $S = S_M$, and start with $n_0 = n \in I_M$, $m_0 = M$, and elapsed time $H_0 = 0$. If $n_0 \notin W_{\eta_{M,M}}$, stop with an initial high failure. Otherwise, while a certified checkpoint $n_i \in I_{m_i}$ has $m_i \ge S$, set

$$q_i = \lfloor r_{hi} m_i \rfloor, \qquad h_i = \tau_{2^{q_i}}(n_i), \qquad n_{i+1} = T^{h_i}(n_i), \qquad H_{i+1} = H_i + h_i. \tag{6.4}$$

If the landing is the upper endpoint 2^{q_i} and $q_i \le S$, follow its deterministic halving orbit to the terminal rank and stop successfully; this costs at most S further steps and cannot create a timeout. If the landing is 2^{q_i} with $q_i > S$, stop with a high endpoint failure. Every nonendpoint landing lies in the unique shell I_{q_i-1}. If $q_i - 1 < S$, hand it to the timeout rule defined below; if $q_i - 1 \ge S$, test it against $W_{\eta_{M,q_i-1}}$. Its first failed test is a high failure, and a successful landing becomes the next checkpoint. As before, a failed landing is never reused.

Lemma 6.1 (possible passage times at large ranks)

There is a constant $K > 0$, depending only on the fixed parameters, such that, for every outer shell I_M and every threshold rank q reached by the high stopped chain, all cumulative first-passage times to its landing band belong to a finite set $H^{\mathrm{hi}}_{M,q}$ satisfying

$$\# H^{\mathrm{hi}}_{M,q} \le K\sqrt{(M+2)\log(M+2)}. \tag{6.5}$$

Proof

Write $g = 1 - a_0$. The point is not that the number of high-rank histories is small. For one block from rank m to threshold rank q, the duration is centered at $(m - q)/g$, with an error of order $\sqrt{m\log(M+2)}$. Every continuing landing lies in the shell I_{q-1}, so the next parent rank is exactly $q - 1$. Consequently the centers telescope along any history to a quantity depending only on the outer rank M and final rank q, while geometric rank decay keeps the accumulated error within $O(\sqrt{M\log(M+2)})$. Thus all cumulative times for fixed M, q lie in one corridor of that width.

Consider one high certified block from a parent shell I_m to threshold 2^q, and let h be its duration and $t = \eta_{M,m}$. The two deterministic orbit envelopes and the first-passage band give

$$(1-t)m - q \le gh < (1+t)(m+1) - q + 1. \tag{6.6}$$

Thus

$$|gh - (m-q)| \le tm + t + 2. \tag{6.6a}$$

The next computation makes this telescoping explicit.

Suppose high blocks $0, \ldots, j$ have been executed. Every landing before the last is a continuing checkpoint in $(2^{q_i-1}, 2^{q_i}]$. By the stopping rule in (6.4), a landing equal to 2^{q_i} terminates the chain: it is followed by deterministic halving when $q_i \le S$, and is a high endpoint failure when $q_i > S$. Thus no continuing checkpoint is an upper endpoint. Each lies in the unique shell I_{q_i-1}, and therefore

$$m_{i+1} = q_i - 1 \qquad (0 \le i < j). \tag{6.6b}$$

This removes the apparent rank branching. In particular,

$$\sum_{i=0}^{j} (m_i - q_i) = M - q_j - j. \tag{6.6c}$$

Summing (6.6a), using (6.6c), and paying one additional unit per block to change the center from $M - q_j - j$ to $M + 1 - q_j$, gives the exact cumulative corridor

$$|gH_{j+1} - ((M+1) - q_j)| \le \sum_{i=0}^{j} (t_i m_i + t_i + 3). \tag{6.6d}$$

For every block,

$$t_i m_i \le D_{hi}\sqrt{m_i \log(M+2)}, \qquad t_i \le \tau,$$

and (6.6b) gives $m_{i+1} \le r_{hi} m_i$. The resulting geometric square-root sum is $O(\sqrt{M\log(M+2)})$, so the right side of (6.6d) has the same order.

For fixed M and $q_j = q$, the center in (6.6d) is fixed. Consequently every feasible integer H_{j+1} lies in one interval of length $O(\sqrt{M\log(M+2)})$; counting its integer points proves (6.5). No union over intermediate high-rank histories remains. No optimality is claimed for this support bound or for its logarithmic factor. □

The fixed-time, fixed-landing count also has the following restricted-time form. If H is any finite set of positive times, $B \subseteq J_Y$, and $D_q \ge 0$ satisfies $D_q/Y \le 1/3$, then

$$\#\left\{n \in I_M : \begin{array}{l}\text{the first passage below } Y \text{ occurs at some } h \in H,\\ T^h(n) \in B,\ E_Y(n) \le D_q\end{array}\right\} \le \#H\left(1 + 3D_q\frac{2^M}{Y}\right)\#B. \tag{6.7}$$

This is Proposition 4.3 summed only over the declared times. No interval structure or density hypothesis on H is used.

Proposition 6.2 (density of large-rank failures)

Assume

$$\frac{D_{hi}}{\sqrt{C_{sw}}} \le \tau \qquad \text{and} \qquad 1 \le S_M \le M.$$

Put

$$c_0 = \frac{1}{2(\log_2 3)^2}, \qquad C_0 = 2e^{4c_0}.$$

For a high threshold rank q, define the complete landing-band target

$$B^{sh}_{M,S_M,q} = \{2^q\} \cup (I_{q-1} \setminus W_{\eta_{M,q-1}}).$$

Then the cap in (6.3) is inactive at every $m \ge S_M$, and

$$\eta^2_{M,m} m = D^2_{hi}\log(M+2). \tag{6.8a}$$

Moreover,

$$\frac{\#(I_M \setminus W_{\eta_{M,M}})}{2^M} \le C_0(M+2)^{-c_0 D^2_{hi}}, \tag{6.8b}$$

and, for every q with $S_M \le q-1$,

$$\frac{|B^{sh}_{M,S_M,q}|}{2^q} \le 2^{-q} + \frac{C_0}{2}(M+2)^{-c_0 D^2_{hi}}. \tag{6.8c}$$

Proof

From $m \ge S_M \ge C_{sw}\log(M+2)$,

$$D_{hi}\sqrt{\frac{\log(M+2)}{m}} \le \frac{D_{hi}}{\sqrt{C_{sw}}} \le \tau.$$

This proves that the cap is inactive and gives (6.8a).

For completeness, the uniform quadratic estimate used here is

$$\frac{\#(I_m \setminus W_t)}{2^m} \le C_0 e^{-c_0 t^2 m} \qquad (0 < t \le 1). \tag{6.8d}$$

When $m \ge 4$ and $tm \ge 2$, apply Lemma 3.1 at height

$$h = \frac{tm}{2\log_2 3}, \qquad 3^h = 2^{tm/2}.$$

Here $h < m/2$, since $t \le 1 < \log_2 3$, so the height is in the range of Lemma 3.1. If $H_m \le h$, then $|d_k| \le h$ for every $k \le m$. Since $n \ge 2^m$,

$$3^{-h} = 2^{-tm/2} \ge n^{-t}, \qquad 3^h = 2^{tm/2} \le 2^{-tm/2} n^t.$$

Thus the multiplicative term in (3.6) lies above $\varrho^k n^{1-t}$ and below $2^{-tm/2}\varrho^k n^{1+t}$. The unused upper-envelope slack is at least

$$(1 - 2^{-tm/2})\,\varrho^m 2^{(1+t)m} \ge \frac{1}{2} 2^{(a_0+t)m}.$$

Meanwhile Lemma 3.2 and $d_k \le h$ bound the additive term in (3.6) by

$$(2+\sqrt{3})3^{2h} = (2+\sqrt{3})2^{tm}.$$

For $m \ge 4$, this is at most the displayed slack because $2^{a_0 m}/2 \ge 2^{4a_0}/2 = 9/2 > 2+\sqrt{3}$. Hence $H_m \le h$ implies $x \in W_t$, or equivalently $I_m \setminus W_t \subseteq \{x \in I_m : H_m(x) > h\}$. By Proposition 2.2, Lemma 3.1, and $I(u) \ge 2u^2$, the complementary proportion is at most

$$2\exp\Big(-\frac{t^2 m}{2(\log_2 3)^2}\Big) = 2e^{-c_0 t^2 m}.$$

Outside this startup regime, $t^2 m \le 4$, so the trivial shell bound is absorbed by $C_0 = 2e^{4c_0}$. Thus (6.8d) is uniform in t. Applying it with $t = \eta_{M,M}$ and using (6.8a) proves (6.8b).

For $q \ge 1$, the landing band lies in I_{q-1} apart from its single upper endpoint 2^q. Apply (6.8d) in I_{q-1}, divide by 2^q, and use (6.8a) at $m = q-1$. This gives (6.8c). □

Consequently the initial high failure and every high-rank landing target have a common density bound

$$d_{hi}(M) \ll (M+2)^{-p_{hi}}, \qquad p_{hi} = \min\{C_{sw}\log 2, c_0 D_{hi}^2\}, \tag{6.9}$$

where $c_0 > 0$ is the fixed constant in the quadratic prefix-deviation bound. The first term in the minimum pays for the dyadic endpoint boundary at the switch.

Critical timeout estimate at small shell ranks

The high construction ends when its landing enters below S_M. At these smaller shell ranks, the proof no longer controls every parity prefix. It replaces that positive all-prefix certificate by a terminal failure test: if the orbit has not crossed its next threshold by the allowed time, then its final odd-step count must be unusually large. That count has the exact binomial law of Proposition 2.2, so its tail can be estimated directly.

Accordingly, declare a low block successful exactly when it reaches its next threshold before its parent-shell rank expires.

Let $L = L_M \asymp \log M$, fix $K_0 > 0$, and put

$$r_L = 1 - \frac{K_0}{2L}, \qquad q_L(m) = \lfloor r_L m \rfloor. \tag{6.10}$$

For sufficiently large L, one has $0 < r_L < 1$ and $q_L(m) < m$ whenever $m \geq L$. For $x \in I_m$, define the low timeout event without assuming eventual hitting:

$$\mathrm{To}_{L,m}(x) \iff T^j(x) > 2^{q_L(m)} \quad \text{for every } 0 \leq j \leq m. \tag{6.11}$$

If this event fails, the least $0 \leq h \leq m$ with $T^h(x) \leq 2^{q_L(m)}$ is an actual successful first-passage block.

When the high stopped chain above lands in a shell below $S = S_M$, admit that landing to the low rule without another all-prefix test. At a low checkpoint $x_i \in I_{m_i}$, stop successfully if $m_i < L$. Otherwise test $\mathrm{To}_{L,m_i}(x_i)$; stop at its first occurrence, and if it does not occur execute the least successful passage just described. If the new threshold rank is below L, stop successfully; otherwise continue from the physical landing. Every successful low transition has duration at most m_i and next shell rank at most $q_L(m_i) < m_i$. Thus the low run is total, closes on its own output, and terminates after at most $S + 1$ stages. Let $\mathrm{Fail}^{\mathrm{to}}_{M,L}$ be the union of the unchanged initial and first high failures with the unique first low timeout.

Lemma 6.3 (timeout targets and feasible times)

Fix $C > 1$. Uniformly for $L \leq m \leq CL$, after one fixed startup,

$$\boxed{\frac{\#\{x \in I_m : \mathrm{To}_{L,m}(x)\}}{2^m} \ll_{C,K_0} m^{-1/2} 2^{-\varkappa_* m}}. \tag{6.12}$$

Moreover, for every sufficiently large outer rank M and every possible first-timeout landing band p, the cumulative times at which that band is reached lie in a finite set $H^{\mathrm{to}}_{M,p}$ satisfying

$$\# H^{\mathrm{to}}_{M,p} \ll \sqrt{(M+2)\log(M+2)}. \tag{6.13}$$

Proof

Let $s = s_m(x)$. The exact affine iterate (2.6) gives

$$T^m(x) < 2 \cdot 3^s + 3^s = 3^{s+1},$$

because $x < 2^{m+1}$ and the additive sum is less than 3^s. Indeed, if the odd letters occur at positions $i_1 < \cdots < i_s$, then $i_j \leq m - s + j - 1$, so that sum is at most

$$\sum_{j=1}^{s} \frac{3^{s-j}}{2^{s-j+1}} = \left(\frac{3}{2}\right)^s - 1 < 3^s$$

(with value zero when $s = 0$). A timeout has $T^m(x) > 2^{q_L(m)}$, and therefore

$$s > q_L(m)\log_3 2 - 1. \tag{6.14}$$

By Proposition 2.2, s_m has the exact $\mathrm{Bin}(m, 1/2)$ counting law on I_m. Put $\theta_{L,m} = r_L m - q_L(m) \in [0,1)$. The normalized real threshold in (6.14) satisfies the exact identity

$$\pi_{L,m} := \frac{q_L(m)p_* - 1}{m} = p_* - \frac{K_0 p_*}{2L} - \frac{\theta_{L,m} p_* + 1}{m} < p_*. \tag{6.14a}$$

Since s is integral, (6.14) is equivalent to

$$s \geq k_{L,m}, \qquad k_{L,m} := \lfloor q_L(m)p_* \rfloor.$$

Put $u_{L,m} = k_{L,m}/m$. Then

$$\pi_{L,m} < u_{L,m} \leq \pi_{L,m} + \frac{1}{m} = \frac{q_L(m)p_*}{m} < p_*,$$

where the last inequality uses $q_L(m) < m$. Thus the timeout cut is subcritical, but differs from p_* by only $O_{C,K_0}(L^{-1})$ when $L \leq m \leq CL$. These thresholds eventually lie in one compact subinterval of $(1/2, 1)$. The ratio of consecutive terms in the upper binomial tail is then bounded strictly below one, so the tail is at most a fixed multiple of its first term. Applying the compact-uniform Stirling estimate from the proof of Lemma 3.1, together with the bounded derivative of $D(p \| 1/2)$ on that interval, gives

$$2^{-m} \sum_{s \geq u_{L,m} m} \binom{m}{s} \ll m^{-1/2} \exp\{-mD(u_{L,m} \| 1/2)\} \ll m^{-1/2} 2^{-\varkappa_* m},$$

because $D(p_* \| 1/2) = \varkappa_* \log 2$. More explicitly, if K_D bounds $|\partial_p D(p \| 1/2)|$ on the fixed compact interval above, then

$$0 \leq m(D(p_* \| 1/2) - D(u_{L,m} \| 1/2)) \leq K_D \left(\frac{CK_0 p_*}{2} + p_* + 1 \right) =: K_{\mathrm{rate}}(C, K_0).$$

Consequently the subcritical cut costs only the fixed multiplier $C_{\mathrm{to}}(C, K_0) := e^{K_{\mathrm{rate}}(C,K_0)}$, which is included in $\ll_{C,K_0}$. In Proposition 6.4, the chosen schedule fixes C and K_0, so that proposition may absorb $C_{\mathrm{to}}(C, K_0)$ without changing the rate $\varkappa_*$. This proves (6.12), including the harmless integer rounding of the first tail index.

It remains to identify the target set and its possible passage times. Suppose the first timeout occurs at a low checkpoint x_i. That checkpoint is the landing of the preceding successful block; let p be that block's threshold rank. The generic form recorded after Lemma 5.1 collapses the preceding threshold chain, making H_i the direct first-passage time of the original source below 2^p, with

$$x_i \in (2^{p-1}, 2^p].$$

The upper endpoint cannot time out, since a power of two crosses every lower dyadic threshold by repeated halving. Hence $x_i \in I_{p-1}$. Since the run would already have stopped below rank L, and timeouts occur only below the switch,

$$L + 1 \leq p \leq S.$$

Thus the first-timeout target in the landing band is

$$\mathcal{C}^{\mathrm{to}}_{L,p} = \{y \in (2^{p-1}, 2^p) \cap \mathbb{N} : \mathrm{To}_{L,p-1}(y)\}, \qquad \frac{|\mathcal{C}^{\mathrm{to}}_{L,p}|}{2^p} \ll p^{-1/2} 2^{-\varkappa_* p}. \tag{6.15}$$

The last estimate is (6.12) with $m = p - 1$, after changing only its fixed constant.

For a fixed high entry threshold, Lemma 6.1 places the cumulative high entry times in an interval of length $O(\sqrt{M \log(M + 2)})$. The possible entry thresholds vary over only $O(S)$ ranks, shifting the common center by $O(S)$. Low ranks decrease strictly, there are at most $S + 1$ successful low stages, and each duration is at most S;

hence

$$\sum_{\text{successful low stages}} h_i \le S(S+1). \tag{6.16}$$

Taking the Minkowski sum enlarges the high interval by $O(S^2)$. Since $S = O(\log(M+2))$, this is $o(\sqrt{M \log M})$, proving (6.13). □

Proposition 6.4 (critical shell failure bound)

Let $L_M \asymp \log M$ be an integer terminal-rank sequence. For every fixed $c > c_*$ and $\beta > 0$, the high parameters and switch constant may be chosen so that the timeout low chain above has, for some fixed $\varepsilon > 0$,

$$\boxed{\frac{\# \operatorname{Fail}^{\mathrm{to}}_{M,L_M}}{2^M} \ll \sqrt{M \log M}\, L_M^{1/2} 2^{-\varkappa_* L_M} + M^{-\varepsilon}.} \tag{6.17}$$

Outside this failure set, some iterate before $c \log n$ lies below 2^{L_M}, and every iterate through the same witness is at most $n^{1+\beta}$.

Proof

Choose

$$a_0 < r_{\mathrm{hi}} < 1 - \frac{1}{c \log 2}, \qquad 0 < \tau < \min\{r_{\mathrm{hi}} - a_0, \beta, a_0\}. \tag{6.18}$$

Because $L_M = O(\log M)$, choose C_{sw} so that $L_M < S_M$ eventually. Increase D_{hi} and then C_{sw}, preserving $D_{\mathrm{hi}}/\sqrt{C_{\mathrm{sw}}} \le \tau$, until the initial and transported high failures total $O(M^{-\varepsilon})$ with a fixed positive margin. Indeed, Proposition 6.2, the direct-passage identity, (6.5), and (6.7) bound their total proportion by

$$d_{\mathrm{hi}}(M) + O(\sqrt{M \log M}\, M^2 d_{\mathrm{hi}}(M)),$$

where the transported contribution is estimated as follows. For a high target rank q, let B_q be its bad landing target. Every producing block has duration at most its parent rank m and $q + 1 > r_{\mathrm{hi}} m$, so the generic rank-scaled-loss statement after Lemma 5.2 gives $E_{2^q}(n) < D_q := (q+2)/r_{\mathrm{hi}}$. Increase the fixed startup so that $D_q/2^q = (q+2)/(r_{\mathrm{hi}} 2^q) \le 1/3$ for every high target rank. Thus (6.7) applies. Since $|B_q| \le 2^q d_{\mathrm{hi}}(M)$, it gives, after division by 2^M, a band proportion $\ll \sqrt{M \log M}(2^{q-M} + D_q) d_{\mathrm{hi}}(M) \ll \sqrt{M \log M}\, M\, d_{\mathrm{hi}}(M)$. Summing over at most $M+1$ possible ranks proves the displayed total. The exponent in $d_{\mathrm{hi}}(M)$ may be made arbitrarily large through the displayed choice of D_{hi} and C_{sw}.

Choose a fixed rational $0 < r_* < r_{\mathrm{hi}}$. Eventually $r_* < r_{L_M}$. Every successful high or low block from parent rank m to threshold rank q then has duration at most m and $q + 1 > r_* m$. The generic rank-chain form recorded after Lemma 5.2 therefore gives, at a first-timeout target rank p,

$$E_{2^p}(n) < \frac{p+2}{r_*}.$$

Since $L_M \asymp \log M$ and $S_M = C_{\mathrm{sw}} \log(M+2) + O(1)$, a fixed $C > 1$, depending only on the chosen schedule, satisfies $p - 1 \le S_M - 1 \le C L_M$ eventually. Thus Lemma 6.3 applies uniformly to every timeout band $L_M + 1 \le p \le S_M$. After increasing the terminal startup, this loss satisfies the smallness condition (4.7). Apply the restricted-time form (6.7) with target $C^{\mathrm{to}}_{L_M,p}$ and time set $H^{\mathrm{to}}_{M,p}$. Equations (6.13) and (6.15) give

$$\frac{\#\{\text{first timeout in band } p\}}{2^M} \ll \sqrt{M \log M}\,(p+1)p^{-1/2}2^{-\varkappa_* p}.$$

The geometric tail satisfies

$$\sum_{p=L_M+1}^{S_M} (p+1)p^{-1/2}2^{-\varkappa_* p} \ll L_M^{1/2}2^{-\varkappa_* L_M}. \tag{6.19}$$

Combining this with the high contribution proves (6.17).

The high-phase clock is at most $M/(1-r_{hi}) + o(M)$, while (6.16) makes the low cost $O((\log M)^2) = o(M)$; a deterministic switch-endpoint tail costs only another $O(S_M)$. The choice (6.18) therefore gives total time below $cM \log 2 \le c \log n$. High blocks use tolerance at most $\tau < \beta$. Every low block starts below 2^{S_M+1} and lasts at most its parent rank. Every state before its final landing is above a threshold greater than one, so its next iterate satisfies $T(z) < 2z$. Consequently each low segment remains below $2^{2S_M+1} = M^{O(1)} = n^{o(1)}$. Together with the high-block envelope and $\tau < \beta$, this proves the stated same-witness ceiling. □

For the main theorem, define the exponent margin

$$\Delta_M = \varkappa_* L_M - \frac{1}{2}\log_2(M+2) - \log_2 \log(M+3). \tag{6.20}$$

If $L_M \asymp \log M$, then the first term in (6.17) is $O(2^{-\Delta_M})$. Thus

$$\boxed{\frac{\#\,\mathrm{Fail}^{\mathrm{to}}_{M,L_M}}{2^M} \ll 2^{-\Delta_M} + M^{-\varepsilon}.} \tag{6.21}$$

Proof of Theorem 1.1

For the bounded sequence (A_M), use the terminal rank in (1.3). Condition (1.4) implies $L_M \asymp \log M$: the upper bound follows from boundedness, and the lower bound follows from the definition of Δ_M. Proposition 6.4 and (6.21) therefore give the shellwise estimate (1.5).

The same condition makes A_M eventually positive. Since the sequence is bounded and $n \in I_M$,

$$2^{L_M} \le 2(M+2)^{A_M} \ll (\log n)^{A_M}, \tag{6.22}$$

with a constant uniform in M. This converts the terminal rank to the stated landing. Finally, the right side of (1.5) tends to zero, so the qualitative part of Lemma 2.1 assembles the shell good sets into one natural-density-one set. □

Proof of Corollary 1.2

For fixed $A > A_{FP}$, take $A_M = A$. Then

$$2^{-\Delta_M} \ll M^{-\varkappa_*(A-A_{FP})}\log M.$$

The high exponent in Proposition 6.4 can be chosen larger than any prescribed $\gamma < \varkappa_*(A - A_{FP})$. Splitting the dyadic shell sum at half the top rank gives the first assertion.

The remaining assertions use Proposition 6.4 directly with the terminal ranks displayed below. Exponentiating 2^{L_M} and using $M \asymp \log n$ then gives their stated landing targets. For the second assertion take

$$L_M = \lceil A_{FP}\log_2(M+2) + B\log_2\log(M+3)\rceil.$$

Then $2^{-\Delta_M} \ll (\log M)^{1-B\varkappa_*}$. The same dyadic split gives every $\gamma < B\varkappa_* - 1$.

For the third assertion take

$$L_M = \lceil A_{FP}\log_2(M+2) + \frac{1}{\varkappa_*}\log_2\log(M+3) + D\log_2\log\log(M+4)\rceil. \tag{6.23}$$

Then $2^{-\Delta_M} \ll (\log\log M)^{-D\varkappa_*}$, and the shell split gives every $\gamma < D\varkappa_*$.

For the final functional statement, let $\Omega(x) \to \infty$. Choose numbers $x_j \uparrow \infty$ so rapidly that $\Omega(x) \geq j$ for every $x \geq x_j$ and $\log j/\log x_j \to 0$. The step function equal to j on $[x_j, x_{j+1})$, with a fixed value on the finite initial interval, is an eventually nondecreasing function satisfying, eventually, $1 \leq \tilde{\Omega} \leq \Omega$ and

$$\tilde{\Omega}(x) \to \infty, \qquad \frac{\log\tilde{\Omega}(x)}{\log x} \longrightarrow 0.$$

Thus $\tilde{\Omega}(x) = x^{o(1)}$. For sufficiently large M, put $z_M = \log(M\log 2)$, so that $z_M \geq 1$ and $z_M \leq \log\log n$ for every $n \in I_M$. Use the terminal rank in (6.23) with the final term replaced by $\log_2\tilde{\Omega}(z_M)$. Since $\log_2\tilde{\Omega}(z_M) = o(\log\log M)$, this rank is still asymptotic to $A_{FP}\log_2 M$, as required in Proposition 6.4, while

$$\Delta_M \geq \varkappa_*\log_2\tilde{\Omega}(z_M) \longrightarrow +\infty.$$

Monotonicity gives an explicitly constructed $\tilde{\Omega}(\log\log n)$-target, which is $o(n)$. Since $\tilde{\Omega} \leq \Omega$ eventually, it also implies the asserted larger Ω-target. □

7. Quantitative companions

The two terminal estimates have different quantitative consequences. Proposition 6.4 gives every strict fixed polylogarithmic exponent and the shell-dependent exponent in the headline. The fixed-barrier failure bound gives the sharper stretched-logarithmic exceptional rate and the clock consequences recorded here.

Proof of Theorem 1.3

We prove the target in (1.8) by choosing the terminal rank so that $2^{L_M} \leq \exp((\log n)^{1-\delta})$ on the source shell.

Fix $0 < \delta < 1$, put $\alpha = 1 - \delta$, and choose a fixed-barrier pair (r, η) satisfying (5.1), with

$$r < 1 - \frac{1}{c\log 2}, \qquad \eta < \beta. \tag{7.1}$$

Take

$$L_M = \lfloor \frac{(M\log 2)^\alpha}{\log 2} \rfloor. \tag{7.2}$$

The factor $M(L_M+1)$ in (5.11) is absorbed by the exponential tail, so

$$\#\ \mathrm{Fail}_{M,L_M} \leq C2^M e^{-c'M^\alpha}. \tag{7.3}$$

For $n \in I_M$,

$$2^{L_M} \leq \exp((M\log 2)^\alpha) \leq \exp((\log n)^\alpha).$$

The clock and path ceiling follow from (5.12) and (7.1). Lemma 2.1, with $\sigma = \alpha$, proves (1.8). □

Raw-clock conversion

Let $s_h(x)$ count the odd states among $x, T(x), \ldots, T^{h-1}(x)$. Direct induction gives

$$\mathrm{Col}^{h+s_h(x)}(x) = T^h(x). \tag{7.4}$$

If $x \in W_\eta$ and $h \le \lfloor \log_2 x \rfloor$, the positive main term in the affine iterate and the upper envelope in (3.8) give

$$s_h(x) \le \frac{h}{2} + \frac{\eta \log x}{\log 3}. \tag{7.5}$$

Thus a fixed-barrier chain has raw time at most

$$\frac{1}{1-r}\left(\frac{3}{2\log 2} + \frac{\eta}{\log 3}\right) \log n + o(\log n). \tag{7.6}$$

For the timeout endpoint, (7.6) applies to the high phase and the low phase has $O(S_M^2) = o(\log n)$ shortcut steps by (6.16); those low steps incur at most twice as many raw steps. As $r \downarrow a_0$ and $\eta \downarrow 0$, the coefficient in (7.6) tends to

$$\frac{3}{\log(4/3)}.$$

The high parameters in (6.18) can therefore be chosen to satisfy any fixed shortcut and raw budgets strictly above their respective limits. To retain the literal raw-orbit ceiling in Corollary 1.4, run the argument with an internal exponent $0 < \beta' < \beta$. Every additional raw odd image is $3m + 1 = 2T(m)$, so $2n^{1+\beta'} \le n^{1+\beta}$ eventually. This proves Corollary 1.4. □

Fixed-power exceptional count

Fix $\alpha > 0$ and $0 < \sigma < 1$. Choose $0 < \delta < 1 - \sigma$. Then

$$\exp((\log n)^{1-\delta}) \le n^{\alpha}$$

eventually, while $1 - \delta > \sigma$. Theorem 1.3 therefore implies (1.10), after decreasing the positive exponential constant and absorbing the finite startup.

The independent fixed-depth argument for the graded clock (1.11) is given in Appendix A. It is not an input to either headline theorem.

8. Scope

The argument proves an orbit-minimum statement on a set of natural density one. It does not prove descent for every starting value, exclude nontrivial cycles, or control the orbit after the selected witness. The endpoint $\delta = 1$ is a strict parameter limit and is not claimed. No finite computation is used as an all-depth premise. The Collatz conjecture predicts $T_{\min}(n) = 1$ for every positive starting value; that universal conclusion is not proved here.

At the polylogarithmic endpoint, the proof requires $\Delta_M \to \infty$. It therefore does not prove the pure target $C(\log n)^{A_{FP}}$, nor the critical secondary target with a bounded tertiary multiplier. The moving theorem improves the terminal scale while weakening its displayed exceptional-set rate; these are separate comparison axes, as recorded in Corollary 1.2.

Theorem 1.1 uses the parity-vector bijection, large-rank prefix-deviation estimate, small-rank timeout tail, first-passage reversal, exact reverse-product loss, direct passage along decreasing thresholds, and the reduced set of possible passage times. Theorem 1.3 instead uses the fixed-barrier failure bound of Theorem 5.3 after the shared first-passage and reverse-loss lemmas; it does not use the small-rank timeout substitution. Neither proof uses a fixed-time endpoint-fiber moment, Fourier estimate, stochastic independence between orbit blocks, mixing statement, Diophantine reduction, or generated-target equidistribution hypothesis. The Boolean walk in Lemma 3.1 is exact counting over the parity words supplied by Proposition 2.2, not a stochastic model of an individual Collatz orbit. The independent graded-clock companion uses only a fixed finite dense-set pullback and is not on either headline dependency chain.

9. Method ceiling and future directions

The endpoint in this paper is set by the first term of (6.17). To separate the leading power from logarithmic factors, suppose a future replacement for the feasible-time summation contributed $M^{\theta+o(1)}$ in place of the present $M^{1/2+o(1)}$, while retaining the same timeout tail and reverse-loss estimate. The scalar budget would then require

$$\varkappa_* L_M > \theta \log_2 M + o(\log M).$$

For $L_M \sim A \log_2 M$, its corresponding method threshold would be $A > \theta/\varkappa_*$. This is only a conditional calculation: the paper proves $\theta = 1/2$, not any smaller value.

Lemma 6.1 already replaces a union over high-rank histories by one interval of feasible times. Further improvement must control the aggregate over the points of that interval, rather than merely reorganize the histories that produced it. Two possible mechanisms are a signed or target-weighted estimate for the actual first-passage landings before taking an absolute upper bound, or a restart-closed renewal estimate for the mass surviving a low-rank timeout. Neither estimate is proved here.

Even eliminating every positive power of M from the feasible-time cost would leave the local factor $L_M^{1/2} 2^{-\varkappa_* L_M}$. It would therefore make the scalar budget compatible with terminal ranks $L_M \to \infty$ arbitrarily slowly, but not with a fixed terminal rank. A fixed bound requires a different survivor-killing estimate whose exceptional mass decays with the outer rank M, for a fixed clock strictly above c_*. If such an estimate reached a fixed interval $[1, C]$, and the finitely many integers in that interval were verified to reach 1, it would imply convergence to 1 on a set of natural density one. That substantially stronger conclusion is outside the present theorem.

Research and software disclosure

Generative AI (including OpenAI Codex) assisted with exploration, formalization workflows, diagnostics, and drafting. The author selected the architecture, checked the mathematics, and accepts responsibility for all content. Finite diagnostics are not premises of any theorem.

The separate Lean package kernel-checks the canonical timeout route to the exported `Main` declaration corresponding to Theorem 1.1. For every bounded exponent sequence with $\Delta_M \to +\infty$, that declaration exposes the literal shell exceptional ratio, moving landing, logarithmic clock, and orbit-height bound through the same witness. The package also exports the fixed-A specialization and the fixed-depth graded-clock companion. Only the canonical timeout route is part of the manuscript-facing formalization surface.

The Lean development has no gaps or admitted obligations: the canonical library builds without `sorry`, `admit`, a project axiom, or a missing module, and the public-root audits report only Lean's standard logical axioms. Some quantitative companions are represented by separate formal projections rather than one declaration with every paper quantifier. In particular, Lean checks the explicit $6.953 \log n$ stretched-logarithmic witness, every strict subendpoint exceptional power, and the raw $10.44 \log n$ stretched-logarithmic landing. The exact endpoint and combined forms in Theorem 1.3, the moving raw specialization in Corollary 1.4, and the literal capstone form of Theorem 5.3 are manuscript derivations assembled from formalized structural components rather than additional public wrapper theorems. This is a declaration-level coverage boundary, not an admitted obligation inside the Lean development. The exact declaration map is recorded in `lean/FORMALIZATION.md`.

The GitHub repository contains the exported theorem. The audited Lean source is frozen at tag `lean-v3.2.0` and commit `ef3410843bf58d69f771f5ba2c0571d54b54da59`. A matching frozen archive is available in the Zenodo software record. The formal package uses Lean `4.15.0` (commit `11651562caae`) and Mathlib revision `9837ca9d65d9de6fad1ef4381750ca688774e608`. From the repository's `lean/` directory, run:

```
lake build
```

The theorem dictionary is `lean/FORMALIZATION.md`. The audit files are `PaperDependencyAudit.lean` for dependencies, `PaperAudit.lean` for the public roots, and `TimeoutEndpointAudit.lean` for the timeout route. The canonical build, placeholder scan, dependency report, and all axiom audits pass, with logical dependencies `propext`, `Classical.choice`, and `Quot.sound`. The default build includes the named canonical audit roots. These checks supplement rather than replace the written proof.

Appendix A. Graded fixed-power clock

This appendix proves the second assertion of Corollary 1.5. The argument is a fixed-depth companion to the headline chain and records a time–descent tradeoff that the stronger polylogarithmic target alone does not express.

Put $g = 1 - a_0$. Fix $a_0 < r < 1$ and $0 < \eta < r - a_0$. For a certified source $x \in I_m$, set $Y_m = 2^{\lfloor rm \rfloor}$, let $G(x)$ be its first entrance into $(Y_m/2, Y_m]$, and put

$$H_m = \left\lceil \frac{(1+\eta-r)m + 2 + \eta}{g} \right\rceil. \tag{A.1}$$

The strict inequality $\eta < r - a_0$ gives $H_m \le m$ eventually. The upper half of the all-prefix envelope (3.8) then gives

$$T^{H_m}(x) < 2^{-gH_m + (1+\eta)(m+1)} \le 2^{rm-1} \le Y_m. \tag{A.2}$$

Hence the first passage defining G occurs by time H_m, and

$$H_m \le \frac{1+\eta-r}{g} m + O_{r,\eta}(1). \tag{A.3}$$

We record why a fixed number of these stages retains natural density one. After absorbing finitely many startup shells, let $\tilde{W}_\eta$ denote the certified set and put $P(S) = \tilde{W}_\eta \cap G^{-1}(S)$. If S is (C, D_{dens})-dense and $0 < \chi < r$, then $P(S)$ is $(C', \chi D_{\mathrm{dens}})$-dense for every sufficiently small fixed $D_{\mathrm{dens}} > 0$. Indeed, for every first passage of length at most H_m, the loss in (4.5) satisfies $E_{Y_m} \le D_{\mathrm{loss}} := H_m/2$, and $H_m/(2Y_m) \le 1/3$ eventually. Apply Proposition 4.3 with $D = D_{\mathrm{loss}}$ and with $B = (Y_m/2, Y_m] \setminus S$. Since $|B| \le C Y_m^{1-D_{\mathrm{dens}}}$, the transported part of the shell complement is

$$\ll C m^2 \frac{2^m}{Y_m} Y_m^{1-D_{\mathrm{dens}}} \ll C m^2 2^{(1-rD_{\mathrm{dens}})m}. \tag{A.4}$$

For fixed $0 < \chi < r$, the polynomial factor is absorbed by $2^{(r-\chi)D_{\mathrm{dens}}m}$, uniformly at cost $O_{r,\chi}(D_{\mathrm{dens}}^{-2})$. Proposition 3.3 absorbs the certification complement after D_{dens} is restricted below a fixed positive threshold. Dyadic summation proves the claimed density update. Iterating it any fixed number of times therefore still gives a natural-density-one set.

Now fix $0 < \alpha < 1$ and $\varepsilon > 0$. Choose $0 < \alpha' < \alpha$ with $c_*(\alpha - \alpha') < \varepsilon/3$, then choose a fixed $R \geq 1$ and $r = (\alpha')^{1/R} > a_0$. Finally choose $0 < \eta < r - a_0$ so that $c_*(1 - \alpha')\eta/(1 - r) < \varepsilon/3$. On the resulting R-fold pullback set, write $n_0 = n$ and $n_{i+1} = G(n_i)$. Finite startup changes only a constant, so $n_i \leq K_* n^{r^i}$ for a fixed K_*, and consequently $n_R \leq K_* n^{\alpha'} \leq n^{\alpha}$ eventually. Summing (A.3) gives

$$\begin{aligned} k_R &\leq \frac{1+\eta-r}{g}\,\frac{1-r^R}{1-r}\log_2 n + O_{\alpha,\varepsilon}(1) \\ &= c_*(1-\alpha')\left(1+\frac{\eta}{1-r}\right)\log n + O_{\alpha,\varepsilon}(1). \end{aligned} \tag{A.5}$$

The two strict parameter margins leave a final third of the clock slack to absorb the fixed remainder, proving (1.11). □

References

[1] R. E. Crandall, *On the "3x+1" problem*, Math. Comp. **32** (1978), no. 144, 1281–1292, doi:10.1090/S0025-5718-1978-0480321-3.

[2] C. J. Everett, *Iteration of the number-theoretic function* $f(2n) = n$, $f(2n+1) = 3n+2$, Adv. Math. **25** (1977), no. 1, 42–45, doi:10.1016/0001-8708(77)90087-1.

[3] M. Inselmann, *An approximation of the Collatz map and a lower bound for the average total stopping time*, arXiv:2402.03276v3 (2024), doi:10.48550/arXiv.2402.03276.

[4] A. V. Kontorovich and J. C. Lagarias, *Stochastic models for the* $3x+1$ *and* $5x+1$ *problems and related problems*, in *The Ultimate Challenge: The* $3x+1$ *Problem* (J. C. Lagarias, ed.), American Mathematical Society, Providence, RI, 2010, pp. 131–188.

[5] I. Korec, *A density estimate for the* $3x+1$ *problem*, Math. Slovaca 44 (1994), no. 1, 85–89, DML-CZ record.

[6] I. Krasikov and J. C. Lagarias, *Bounds for the* $3x+1$ *problem using difference inequalities*, Acta Arith. **109** (2003), 237–258, doi:10.4064/aa109-3-4.

[7] J. C. Lagarias, *The* $3x+1$ *problem and its generalizations*, Amer. Math. Monthly **92** (1985), no. 1, 3–23, doi:10.2307/2322189.

[8] J. C. Lagarias and A. Weiss, *The* $3x+1$ *problem: two stochastic models*, Ann. Appl. Probab. **2** (1992), no. 1, 229–261, doi:10.1214/aoap/1177005779.

[9] I. A. Shaik, *Quantitative Collatz Descent to Stretched-Logarithmic Scale in Natural Density, with a Lean 4 Formalization*, version 2.0.2, preprint (2026), doi:10.5281/zenodo.21851173.

[10] T. Tao, *Almost all orbits of the Collatz map attain almost bounded values*, Forum Math. Pi **10** (2022), e12, doi:10.1017/fmp.2022.8.

[11] R. Terras, *A stopping time problem on the positive integers*, Acta Arith. **30** (1976), 241–252.

[12] R. Terras, *On the existence of a density*, Acta Arith. **35** (1979), 101–102, doi:10.4064/aa-35-1-101-102.

[13] G. J. Wirsching, *The Dynamical System Generated by the* $3n+1$ *Function*, Lecture Notes in Mathematics, vol. 1681, Springer-Verlag, Berlin, 1998, doi:10.1007/BFb0095985.