Forum of Mathematics, Sigma

CAMBRIDGE
UNIVERSITY PRESS

RESEARCH ARTICLE

Quantitative Collatz Descent to Stretched-Logarithmic Scale in Natural Density

Idris Ali Shaik[1]

[1]Independent researcher; E-mail: shaikidris@apache.org.

Keywords: Collatz map, natural density, Rényi moment, parity vector, endpoint fiber, quantitative descent

MSC Codes: *Primary* – 11B83; *Secondary* – 60F10

Abstract
Let $T(n) = n/2$ for even n and $T(n) = (3n+1)/2$ for odd n, and write $T_{\min}(n) = \min_{k\geq 0} T^k(n)$. For every $0 < \delta < \delta_{\text{end}} = 0.251245530155874\ldots$, we prove that $T_{\min}(n) \leq \exp((\log n)^{1-\delta})$ on a set of natural density one, where $\delta_{\text{end}} = \log(1/a_0)/\log(2/a_0)$ and $a_0 = (\log_2 3)/2$. Earlier published natural-density theorems give fixed-power bounds; the present threshold is smaller than every fixed power. The descent is witnessed within $6.953 \log n$ half-map iterations. Moreover, for every $0 < \sigma < 1 - \delta/\delta_{\text{end}}$, the number of exceptions up to X is at most $5X\exp(-c_{\delta,\sigma}(\log X)^{\sigma})$ for all sufficiently large X. The main input is a central fixed-total Rényi estimate for one fixed $1/2 < \theta < 1$, which yields an asymptotically linear density pullback and an endpoint-only bootstrap. The principal theorem chain is formalized in Lean 4. In a smaller range, a companion theorem gives a shrinking-error two-sided orbit envelope and a uniform quantitative first-passage profile. These almost-all results neither prove descent for every initial value nor exclude exceptional cycles or divergent trajectories.

1. Introduction and main results

1.1. Background and quantitative setting

The Collatz conjecture asserts eventual arrival at 1 for every positive integer. We study a weaker quantitative question: how far one can certify descent for all but a sparse set of initial values, with the sparsity measured in natural density. The exact parity encoding used below originates in the work of Terras [8] and Everett [3]. Allouche [2] proved the natural-density estimate $T_{\min}(n) \leq n^{\theta}$ for every fixed $\theta > 3/2 - (\log 2)/(\log 3)$. Korec [5] extended this to every fixed $\theta > \log_4 3$, while Inselmann [4, Theorem 1.1 and Corollary 1.4] proved $T_{\min}(n) \leq n^{\varepsilon}$ in natural density for every fixed $\varepsilon > 0$, together with a fixed-tolerance two-sided envelope through the horizon $2\log n/\log(4/3)$. For every $0 < \delta < \delta_{\text{end}}$, Theorem 1.1 gives the eventually smaller threshold $T_{\min}(n) \leq \exp((\log n)^{1-\delta})$. Relative to Inselmann's fixed-power natural-density conclusion, the advance is along the explicit target-scale and exceptional-rate axes rather than the limiting iteration-depth axis. The quantitative transport is obtained from the fixed-total endpoint-fiber Rényi bounds of Theorems 3.8 and 3.9, iterated through the endpoint-only bootstrap of Section 6. The full-envelope theorem gives an explicit shrinking tolerance through the same limiting horizon.

Tao [7] proved that, for every function tending to infinity, the initial values whose orbit attains that threshold form a set of logarithmic density one. This gives a much stronger orbit threshold in a different density notion; Section 9 compares the two settings precisely.

The quantitative contribution here is an independently derived explicit stretched-logarithmic rate, its stretched-exponential exceptional count, and the associated half-map witness. The manuscript does not claim priority for arbitrary-diverging-threshold natural-density descent.

All unqualified logarithms are natural; $\log_2$ denotes the base-two logarithm. For $S \subseteq \mathbb{N}_{\geq 1}$, put

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

The set S is (C, D)-dense if

$$B_S(X) \leq CX^{1-D} \qquad (X \geq 1),$$

where $C > 0$ and $0 < D \leq 1$. In particular, every such set has natural density one. This agrees with the $*$-density terminology of [4, Definition 2.6].

For a set S, define its logarithmic-block Collatz pullback by

$$\mathrm{Pull}(S) = \{n \geq 1 : T^{\lfloor \log_2 n \rfloor}(n) \in S\}.$$

The pullback is the density-propagation operation iterated in the envelope bootstrap.

1.2. Main theorem and quantitative refinements

Put

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

Theorem 1.1 has endpoint

$$\delta_{\mathrm{end}} = \frac{\log(1/a_0)}{\log(2/a_0)} = 0.251245530155874\ldots \tag{1.1}$$

and is a nonattained supremum over fixed parameters $\theta < 1$.

Theorem 1.1 (Stretched-logarithmic natural-density descent)**.** *For every* $0 < \delta < \delta_{\mathrm{end}}$,

$$T_{\min}(n) \leq \exp\left((\log n)^{1-\delta}\right) \tag{1.2}$$

on a set of natural density one.

The displayed expression is needed only outside a finite initial interval. Any totalization convention for the finitely many smaller inputs, including $n = 1$, does not affect natural density.

In decimal-digit terms, the theorem supplies a value in the orbit with at most

$$1 + \frac{(\log n)^{1-\delta}}{\log 10}$$

digits. Since n itself has $\asymp \log n$ digits, the certified digit-count ratio tends to zero. This is qualitatively stronger than a fixed-power estimate $T_{\min}(n) \leq n^{\alpha}$, which only reduces the digit count by the constant factor α. This comparison concerns the threshold itself; the density notion and description of the underlying good set remain separate features of each theorem.

The conclusion is an almost-all statement in natural density. It does not assert descent for every initial value, and it is compatible with arbitrarily long individual exceptional trajectories.

Corollary 1.2 (Quantitative exceptional count)**.** *For every* $0 < \delta < \delta_{\mathrm{end}}$ *and every*

$$0 < \sigma < 1 - \frac{\delta}{\delta_{\mathrm{end}}},$$

there are $c_{\delta,\sigma} > 0$ *and* $X_{\delta,\sigma}$ *such that*

$$\#\left\{1 \le n \le X : T_{\min}(n) > \exp\left((\log n)^{1-\delta}\right)\right\} \le 5X \exp\left(-c_{\delta,\sigma}(\log X)^{\sigma}\right) \tag{1.3}$$

for every $X \ge X_{\delta,\sigma}$*. In particular, one may choose* $\sigma > \delta$ *throughout the larger range*

$$0 < \delta < \frac{\delta_{\text{end}}}{1+\delta_{\text{end}}} = 0.2007963458\ldots.$$

The constants $c_{\delta,\sigma}$ *and* $X_{\delta,\sigma}$ *are effective after the fixed interior parameters are chosen; their closed forms are not optimized, and no numerically practical onset scale is claimed. The restriction on* σ *comes from the endpoint-block schedule; the dyadic summation preserves the stretched-exponential power and changes only constants.*

Corollary 1.3 (Bounded witnessing time)**.** *For every* $0 < \delta < \delta_{\text{end}}$*, the density-one set in Theorem 1.1 may be chosen so that every sufficiently large n in it has an integer k, counting half-map iterations, with*

$$0 \le k < 6.953 \log n$$

for which

$$T^k(n) \le \exp\left((\log n)^{1-\delta}\right). \tag{1.4}$$

Moreover, for every fixed $\beta > 0$*, the density-one set may be chosen (depending on* β*) so that the witness k above also satisfies*

$$\max_{0\le j\le k} T^j(n) \le n^{1+\beta}.$$

This ceiling concerns the orbit segment up to the descent witness k; it does not assert control at later times below the numerical bound $6.953 \log n$.

1.3. Companion theorem and further consequences

For the pathwise companion put

$$b = 1 - a_0, \qquad \rho = \frac{\sqrt{3}}{2}, \qquad c_{\text{asym}} = \frac{\log 3}{6\log 2},$$

and

$$\delta_{\text{env}} = \left[3 + \frac{\log\left(1/\left(c_{\text{asym}}(\sqrt{2}-1)^2\right)\right)}{\log(1/a_0)}\right]^{-1} = 0.06134085341\ldots. \tag{1.5}$$

Theorem 1.4 (Full-envelope theorem)**.** *For every* $0 < \delta < \delta_{\text{env}}$*, there are a number* $\upsilon > \delta$ *and a set* $\mathscr{H}$ *of natural density one such that every sufficiently large* $n \in \mathscr{H}$*, with* $M = \lfloor \log_2 n \rfloor$*, admits numbers*

$$\lambda_n, t_n \ll (\log n)^{-\upsilon}$$

for which

$$\rho^k n^{1-t_n} \le T^k(n) \le \rho^k n^{1+t_n}$$

uniformly for

$$0 \le k \le \left\lfloor \frac{(1-\lambda_n)M}{b} \right\rfloor.$$

Moreover,

$$T_{\min}(n) \le \exp\big((\log n)^{1-\delta}\big)$$

on the same set. In particular, the entire certified segment has maximum $n^{1+o(1)}$ and follows the geometric rate $\rho = \sqrt{3}/2$ up to a subpower error.

The same construction controls every intermediate iterate, not only the logarithmic-block endpoints. Its constants are effective after the fixed parameters are chosen, but the resulting starting scale is a proof constant, not an estimate of typical behavior.

Within the proved envelope recurrence, Lemma 8.2 shows that δ_{env} is the exact supremum imposed by $q > a_0$, $\kappa < \sqrt{2} - 1$, and $c_{\mathrm{L}} < c_{\mathrm{asym}}$. Neither this range nor the admissible υ is claimed sharp under a stronger pullback or concatenation theorem.

1.4. *First-passage and fixed-power consequences*

Explicit subpower form.

For every fixed $0 < \delta < \delta_{\mathrm{end}}$, there is a non-increasing function $\varepsilon_\delta : \mathbb{N}_{\ge 1} \to (0, \infty)$, with $\varepsilon_\delta(n) \to 0$ and

$$\varepsilon_\delta(n) = (\log n)^{-\delta}$$

outside a finite initial interval, such that

$$T_{\min}(n) \le n^{\varepsilon_\delta(n)}$$

on a set of natural density one. Indeed, take $\varepsilon_\delta(n) = (\log n)^{-\delta}$ for $n \ge 2$, set $\varepsilon_\delta(1) = \varepsilon_\delta(2)$, and apply Theorem 1.1, using $n^{(\log n)^{-\delta}} = \exp((\log n)^{1-\delta})$ for $n \ge 2$.

Corollary 1.5 (Uniform quantitative first-passage profile)**.** *For a positive threshold function h, write*

$$\tau_h(n) = \min\{k \ge 0 : T^k(n) \le h(n)\} \tag{1.6}$$

whenever this set is nonempty, and put

$$c_* = \frac{2}{\log(4/3)} = 6.95211899\ldots.$$

For every fixed $0 < \delta < \delta_{\mathrm{env}}$, there are exponents $\delta < \upsilon < \sigma < 1$, constants $C_0, c_0, C_1 > 0$, and a set A_δ such that

$$\#([1, X] \setminus A_\delta) \le C_0 X \exp\left(-c_0 (\log X)^{\sigma}\right) \qquad (X \ge 2),$$

and, for every sufficiently large $n \in A_\delta$, the first-passage time is defined simultaneously for every $(\log n)^{-\delta} \le y \le 1$ and satisfies

$$\sup_{(\log n)^{-\delta} \le y \le 1} |\tau_{n^y}(n) - c_*(1 - y)\log n| \le C_1 (\log n)^{1-\upsilon}. \tag{1.7}$$

In particular, for $h_\delta(n) = \exp((\log n)^{1-\delta})$,

$$\tau_{h_\delta}(n) = c_* \log n - c_* (\log n)^{1-\delta} + O\left((\log n)^{1-\upsilon}\right),$$

while every fixed $0 < x < 1$ satisfies, on the same set,

$$\tau_{n^x}(n) = (1 - x) c_* \log n + O\left((\log n)^{1-\upsilon}\right). \tag{1.8}$$

Thus the full-envelope theorem localizes, with one quantitative exceptional set, the passage through every power scale between n and the stretched-logarithmic terminal threshold. In particular, $\tau_{h_\delta}(n)/\log n \to c_*$ and $\tau_{n^x}(n)/\log n \to (1-x)c_*$ along A_δ.

The leading fixed-power law in (1.8) also follows from the fixed-tolerance envelope in [4, Theorem 1.1] by a density-one diagonalization in the tolerance. The uniform error in (1.7), its exceptional count, and the second-order shrinking-target term instead use the quantitative shrinking-tolerance estimate and do not follow from a single fixed-tolerance envelope.

Corollary 1.6 (Power-saving fixed-power descent)**.** *For every*

$$\beta > \delta_{\text{end}}^{-1} = 3.9801703114\ldots,$$

there are constants $K_\beta, c_\beta > 0$ such that, for every fixed $0 < \alpha < 1$, there is a constant $C_{\alpha,\beta} > 0$ satisfying

$$\#\{1 \le n \le X : T_{\min}(n) > K_\beta n^\alpha\} \le C_{\alpha,\beta} X^{1-c_\beta \alpha^\beta} \qquad (X \ge 1). \tag{1.9}$$

The constants K_β, c_β are uniform in α; the density prefactor $C_{\alpha,\beta}$ is not asserted to be uniform as $\alpha \downarrow 0$.

1.5. Method and organization

The critical-moment endpoint-only boundary, used below as a baseline, is

$$\delta_{\text{crit}} = \frac{\log(1/a_0)}{\log(3/a_0)} = 0.174719543136035\ldots. \tag{1.10}$$

The endpoint-only recurrence permits

$$\delta < \frac{\log(1/r)}{\log(1/\chi)}.$$

At the critical Rényi parameter $\theta = 1/2$, taking $r \downarrow a_0$ and $\chi \uparrow a_0/3$ gives δ_{crit}. In the central higher-moment range, $\chi \uparrow \Lambda(\theta) = a_0\theta/(1+\theta)$; then $\theta \uparrow 1$ gives $\chi \uparrow a_0/2$ and hence $\delta_{\text{end}} = \log(1/a_0)/\log(2/a_0)$. These boundaries are nonattained because the parameter choices remain strict. The smaller δ_{env} is the range in which the full two-sided envelope is retained.

The two principal results have the dependency structure

$$\begin{aligned}\text{central fixed-total Rényi estimate} &\Longrightarrow \text{higher-moment pullback}\\ &\Longrightarrow \text{endpoint-only bootstrap} \Longrightarrow \text{Theorem 1.1},\end{aligned}$$

while

$$\text{critical moment} + \text{two-sided envelope concatenation} \Longrightarrow \text{Theorem 1.4}.$$

The proof has ten steps.

1. The Terras parity bijection identifies a dyadic shell with all Boolean parity words of the corresponding length.
2. At a fixed odd count, equality of two endpoints forces congruence of their affine corrections modulo the corresponding power of 3.
3. Fixing the first odd position identifies the remaining parity word with a positive composition of a fixed total.
4. Distinct compositions give distinct integer affine numerators. Hence a modular correction fiber of size r occupies at least r distinct integer quotient levels.
5. A reverse-spend estimate controls every fixed central Rényi parameter $1/2 < \theta < 1$; cohorts outside the fixed central window have exponentially small source mass.

6. A heavy-fiber split gives the nonlinear pullback rate $\theta\psi(D)/(1+\theta)$, hence every small-D transport constant below $a_0\theta/(1+\theta)$.
7. A fixed maximal-barrier window supplies one dense set on which the logarithmic-block endpoint map satisfies $F(n) \ll n^r$, for any fixed $r > a_0$.
8. Iterating only these endpoints through $O(\log\log n)$ blocks and then taking fixed θ arbitrarily close to one gives the boundary $\log(1/a_0)/\log(2/a_0)$ in Theorem 1.1.
9. The endpoint schedule also gives the full-range witnessing-time and the optimized exceptional-count range $\sigma < 1 - \delta/\delta_{\mathrm{end}}$. Retaining the stronger two-sided envelope gives Theorem 1.4.
10. A phase-transition lemma records the critical $\theta = 1/2$ baseline, the fixed-θ endpoints, and their nonattained limit as $\theta \uparrow 1$.

Sections 2–5 establish the preliminaries, fixed-total estimates, and density pullback. Section 6 proves Theorem 1.1, Sections 7–8 develop the full-envelope companion, and Section 9 records scope and literature comparisons.

2. Quantitative density and Collatz preliminaries

For comparison with literature using the unaccelerated map, put

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

Every step of T is either one or two steps of Col, and the extra raw iterate after an odd input is $3n+1 = 2T(n) > T(n)$. Consequently

$$\min_{j\ge 0} \mathrm{Col}^j(n) = T_{\min}(n),$$

and a witness after k half-map steps gives a raw-Collatz witness after some m with $k \le m \le 2k$.

We use two elementary dyadic summation facts. First, if

$$\#\big(S^c \cap [2^M, 2^{M+1})\big) \le K e^{-\gamma M} 2^M \qquad (M \ge 0),$$

where $K > 0$ and $0 < \gamma < \log 2$, then summing the geometric series shows that S is

$$\left(\frac{2K}{2e^{-\gamma}-1}, \frac{\gamma}{\log 2}\right)\text{-dense.}$$

Second, if

$$\#\big(E \cap [2^M, 2^{M+1})\big) \le \delta_M 2^M \qquad \text{with } \delta_M \longrightarrow 0,$$

then E has natural density zero.

Proof. Indeed, after fixing $\epsilon > 0$, all sufficiently late shells contribute at most $\epsilon \sum 2^M$, while the finitely many early shells contribute $O_\epsilon(1)$. Dividing at $2^J \le X < 2^{J+1}$, taking a limsup, and then letting $\epsilon \downarrow 0$ proves the claim. □

The quantitative version used later is equally elementary.

Lemma 2.1 (Varying-rate dyadic summation)**.** *Let $0 < \sigma \le 1$, $A, \gamma > 0$, and suppose that, for all sufficiently large M,*

$$\#\big(E \cap [2^M, 2^{M+1})\big) \le A e^{-\gamma (M+4)^\sigma} 2^M. \tag{2.1}$$

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

$$\#(E \cap [1, X]) \le A' X \exp\left(-\gamma' (\log X)^\sigma\right) \qquad (X \ge X_0). \tag{2.2}$$

Proof. One may take

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

Let M_* be a threshold beyond which (2.1) holds. It is enough to take

$$X_0 \geq \max\{4, 2^{2M_*}\},$$

and, when $0 < \sigma < 1$, also require

$$\log X_0 \geq (2\gamma')^{1/(1-\sigma)}.$$

Indeed, let $2^J \leq X < 2^{J+1}$. Then $\lfloor J/2 \rfloor \geq M_*$. The shells with $M < \lfloor J/2 \rfloor$ contain fewer than

$$\sum_{M < \lfloor J/2 \rfloor} 2^M < 2^{\lfloor J/2 \rfloor} \leq X^{1/2}$$

integers. On every remaining shell,

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

Their contribution is therefore at most

$$A\exp\left(-\gamma'(\log X)^{\sigma}\right)\sum_{M \leq J} 2^M \leq 2AX\exp\left(-\gamma'(\log X)^{\sigma}\right).$$

Finally, the choice of X_0 gives

$$\gamma'(\log X)^{\sigma} \leq \frac{1}{2}\log X.$$

For $\sigma = 1$, the same inequality follows directly from $\gamma' \leq 1/2$. Hence

$$X^{1/2} \leq X\exp\left(-\gamma'(\log X)^{\sigma}\right).$$

Adding the early and late contributions proves (2.2) with the displayed choices of A' and γ'. □

We shall also use the following elementary symmetric-binomial estimate. If $X \sim \mathrm{Bin}(n, 1/2)$ and $y \geq 0$, then

$$\Pr\left(\left|X - \frac{n}{2}\right| \geq y\right) \leq 2\exp\left(-\frac{2y^2}{n}\right) \qquad (n \geq 1).$$

Indeed, for $\lambda > 0$,

$$\mathbb{E}e^{\lambda(X-n/2)} = \left(\cosh\frac{\lambda}{2}\right)^n \leq e^{n\lambda^2/8}.$$

Proof. For completeness, $\cosh u \leq e^{u^2/2}$: the even function $u^2/2 - \log\cosh u$ has derivative $u - \tanh u$ on $[0, \infty)$, whose derivative is $\tanh^2 u \geq 0$, and it vanishes at zero. Markov's inequality with $\lambda = 4y/n$ gives the upper tail, and applying the same argument to $-X$ gives the lower tail. Thus no external concentration theorem is used below. □

2.1. Elementary density calculus

We use the following consequences of the definition. If S_i is (C_i, D_i)-dense for $i = 1, 2$, then

$$S_1 \cap S_2 \quad \text{is} \quad (C_1 + C_2, \min(D_1, D_2))\text{-dense}.$$

Indeed,

$$B_{S_1 \cap S_2}(X) \le B_{S_1}(X) + B_{S_2}(X) \le (C_1 + C_2) X^{1-\min(D_1, D_2)}.$$

Also, if

$$S^c \subseteq S_0^c \cup [1, F]$$

and S_0 is (C, D)-dense, then S is $(C + F, D)$-dense, because

$$B_S(X) \le B_{S_0}(X) + F \le (C + F) X^{1-D}.$$

Finally, decreasing the exponent preserves density: if $0 < D' \le D$ and S is (C, D)-dense, then it is also (C, D')-dense.

We shall also use the following elementary density-one diagonalization. Let $u_n^{(1)}, \dots, u_n^{(d)}$ be finitely many real sequences. If, for every $j \ge 1$, the set on which all $u_n^{(i)}$ are defined and

$$\max_{1 \le i \le d} |u_n^{(i)} - L_i| < \frac{1}{j}$$

has natural density one, then there is a single set A of natural density one such that $u_n^{(i)} \to L_i$ as $n \to \infty$ through A, for every i. Indeed, let B_j be the complementary bad set. The sets B_j increase with j, and each has density zero. Choose strictly increasing N_j so that

$$\#(B_j \cap [1, N]) \le \frac{N}{j} \qquad (N \ge N_j),$$

retain all integers below N_1, and on $N_j \le n < N_{j+1}$ retain exactly the integers outside B_j. For $N_j \le N < N_{j+1}$, every discarded integer up to N belongs to B_j, so the discarded proportion is at most $1/j$. The retained set therefore has density one, and its block-j errors are less than $1/j$.

For $n \ge 1$, define the parity bits and prefix odd counts by

$$p_i(n) = \mathbf{1}_{\{T^i(n) \text{ odd}\}}, \qquad s_k(n) = \sum_{i=0}^{k-1} p_i(n). \tag{2.3}$$

Proposition 2.2 (Terras parity bijection). *For every $M \ge 0$, the map*

$$n \bmod 2^M \longmapsto \big(p_0(n), \dots, p_{M-1}(n)\big) \tag{2.4}$$

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

$$\#\{n \in [2^M, 2^{M+1}) : s_M(n) = s\} = \binom{M}{s}. \tag{2.5}$$

Proof. For completeness, the one-step 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 $n \bmod 2^M$. Conversely, equal first bits give equal parity at time zero and equal parity words after applying T. On the even branch, lifting from $T(n) \bmod 2^k$ is $n = 2T(n) \bmod 2^{k+1}$. On the odd branch,

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

where $2T(n) \bmod 2^{k+1}$ is determined by $T(n) \bmod 2^k$, and 3 is a unit modulo 2^{k+1}. Induction therefore recovers $n \bmod 2^M$. The two finite sets have the same cardinality, proving the bijection. Each residue occurs exactly once in the dyadic shell, and (2.5) follows. □

Define the integer affine correction recursively by

$$c_0(n) = 0, \qquad c_{k+1}(n) = \begin{cases} c_k(n), & p_k(n) = 0, \\ 3c_k(n) + 2^k, & p_k(n) = 1. \end{cases} \tag{2.6}$$

Proposition 2.3 (Exact affine iterate). *For every $n \geq 1$ and $k \geq 0$,*

$$2^k T^k(n) = 3^{s_k(n)} n + c_k(n). \tag{2.7}$$

Moreover,

$$0 \leq c_k(n) < 3^{s_k(n)} 2^k. \tag{2.8}$$

Proof. Both statements follow simultaneously by induction on k. The even step leaves the correction unchanged; the odd step multiplies the previous correction by 3 and adds 2^k. More explicitly, if the k-th letter is even, then

$$c_{k+1} = c_k < 3^{s_k} 2^k < 3^{s_{k+1}} 2^{k+1}.$$

If it is odd, then $s_{k+1} = s_k + 1$, and

$$c_{k+1} = 3c_k + 2^k < 3^{s_k+1} 2^k + 2^k < 3^{s_k+1} 2^{k+1},$$

because $1 < 3^{s_k+1}$. This proves the strict upper bound in (2.8) at the next step and completes the simultaneous induction. □

Corollary 2.4 (Endpoint range). *If $n \in [2^M, 2^{M+1})$ and $s_M(n) = s$, then*

$$0 < T^M(n) < 3^{s+1}. \tag{2.9}$$

Proof. Indeed, (2.7), (2.8), and $n < 2^{M+1}$ give

$$2^M T^M(n) < 3^s 2^{M+1} + 3^s 2^M = 3^{s+1} 2^M.$$

□

3. Fixed-total endpoint coding and quotient order

For the half-Collatz map, let

$$f_{M,s}(y) = \#\{n \in [2^M, 2^{M+1}) : s_M(n) = s,\ T^M(n) = y\},$$

$$\mathcal{Y}_s = 3^{s+1}, \qquad g_{M,s}(y) = \frac{\mathcal{Y}_s f_{M,s}(y)}{\binom{M}{s}},$$

and

$$G_M(\theta) = 2^{-M} \sum_{s,y} f_{M,s}(y) g_{M,s}(y)^{\theta}.$$

The reference size $\mathcal{Y}_s$ is the endpoint-range bound supplied by Corollary 2.4.

Theorem 3.1 (Critical near-L^1 endpoint information)**.** *For every fixed $0 < \theta < 1/2$, there is $C_\theta < \infty$ such that, for every $M \geq 1$,*

$$G_M(\theta) \leq C_\theta (M+1)^{\theta}. \tag{3.1}$$

At the critical endpoint one has the explicit bound

$$G_M\left(\frac{1}{2}\right) \leq \frac{9\sqrt{3}}{4} M. \tag{3.2}$$

Corollary 3.2 (Coarse polynomial majorant)**.** *For every $M \geq 1$,*

$$G_M\left(\frac{1}{2}\right) \leq 9\sqrt{3}\, M^{3/2}. \tag{3.3}$$

This weaker polynomial majorant follows immediately from (3.2), since $M \geq 1$ and

$$\frac{9\sqrt{3}}{4} M \leq 9\sqrt{3}\, M^{3/2}.$$

It leaves every exponential rate in the critical pullback and Theorem 1.4 unchanged. It is not used in the central higher-moment proof of Theorem 1.1.

Corollary 3.3 (Heavy-fiber mass)**.** *For every $\nu > 0$,*

$$2^{-M} \sum_{\substack{s,y \\ g_{M,s}(y) > e^{\nu M}}} f_{M,s}(y) \leq \frac{9\sqrt{3}}{4} M e^{-\nu M/2}. \tag{3.4}$$

This is Markov's inequality applied to (3.2). Thus exponentially overloaded endpoint fibers may exist, but they carry exponentially little source mass. The coarser Corollary 3.2 replaces $(9\sqrt{3}/4)M$ by $9\sqrt{3}\, M^{3/2}$.

Theorem 3.4 (Exact nonlinear arbitrary-target pullback)**.** *Define*

$$\psi(D) = -\log\left(\frac{1+3^{-D}}{2}\right).$$

For every $0 < \eta < 1$, there is $K_\eta > 0$ such that every (C, D)-dense set, $0 < D \leq 1$, has pullback

$$\left(K_\eta (C+1)\psi(D)^{-1}, \frac{\eta\psi(D)}{3\log 2}\right)\text{-dense}. \tag{3.5}$$

The proof uses neither a universal DM-wide parity cone nor subexponential normalized L^2 collision energy.

Corollary 3.5 (Coarse recurrence prefactor)**.** *For every $0 < \eta < 1$, there is $K'_\eta > 0$ such that the pullback in Theorem 3.4 is*

$$\left(K'_\eta (C+1) D^{-2}, \frac{\eta\psi(D)}{3\log 2}\right)\text{-dense}. \tag{3.6}$$

This follows from Theorem 3.4 and the chord bound $\psi(D) \geq D\log(3/2)$; the extra power of D^{-1} is retained only to give a simple uniform recurrence.

Corollary 3.6 (Uniform linear pullback)**.** *Uniformly for* $0 < D \leq 1$*, every*

$$0 < \xi < \frac{\log(3/2)}{3\log 2} \tag{3.7}$$

is admissible as the exponent ξD*, after changing the prefactor in* (3.6)*.*

Corollary 3.7 (Asymptotic small-D pullback)**.** *As* $D \downarrow 0$*, every*

$$0 < \xi < \frac{\log 3}{6\log 2} \tag{3.8}$$

is admissible as the exponent ξD *for all sufficiently small* D*.*

Both endpoints in (3.7)–(3.8) are suprema, not attained constants.

Theorem 3.8 (Central higher-Rényi information)**.** *Fix* $1/2 < \theta < 1$*, and define the available central half-width*

$$\varepsilon_0(\theta) = \frac{2^\theta - 1}{3^\theta - 1} - \frac{1}{2}. \tag{3.9}$$

The inequality $\varepsilon_0(\theta) > 0$ *is proved at the beginning of the proof in Section 5. Fix a number* $0 < \varepsilon < \varepsilon_0(\theta)$*. For a first-odd-position cohort* u*, put* $N = M - u$*, and retain it only when*

$$\left| s - \frac{N+1}{2} \right| \leq \varepsilon N.$$

Let $f^{\mathrm{cen}}_{M,s}(y)$ *be the endpoint multiplicity obtained by summing only those retained cohorts, and define*

$$G^{\mathrm{cen}}_M(\theta,\varepsilon) = 2^{-M} \sum_{s,y} f^{\mathrm{cen}}_{M,s}(y) \left(\frac{3^{s+1} f^{\mathrm{cen}}_{M,s}(y)}{\binom{M}{s}} \right)^{\theta}. \tag{3.10}$$

There are constants $C_{\theta,\varepsilon}, B_{\theta,\varepsilon} < \infty$*, depending only on* θ *and* ε*, such that*

$$G^{\mathrm{cen}}_M(\theta,\varepsilon) \leq C_{\theta,\varepsilon}(M+1)^{B_{\theta,\varepsilon}} \qquad (M \geq 1) \tag{3.11}$$

The requirement that ε be fixed independently of D is essential for the bootstrap. The proof in Section 5 establishes the polynomial bound with $B_{\theta,\varepsilon} = 1 + \theta$, while the symmetric-binomial estimate proved in Section 2 shows that the omitted source mass is $O_\varepsilon(e^{-c_\varepsilon M})$.

Under the uniform positive-composition model, $J_N - 1 \sim \mathrm{Bin}(N-1, 1/2)$, so the retained cohorts form a fixed typical window for the separator density. For $\pi = (s-1)/(N-1)$, the reverse-spend reproduction factor $\mathscr{R}_\theta(\pi)$ is less than one precisely when $\pi < \alpha_\theta = (2^\theta - 1)/(3^\theta - 1)$, as recorded in Proposition 6.3. Thus $\varepsilon < \varepsilon_0(\theta)$ keeps every retained cohort uniformly subcritical, while the omitted cohorts have a fixed exponential source-mass bound. Independence from D lets this discarded-mass estimate remain uniform as the useful transport exponent shrinks through successive pullback stages.

Theorem 3.9 (Central higher-moment pullback)**.** *Fix* $1/2 < \theta < 1$ *and* $0 < \varepsilon < \varepsilon_0(\theta)$*, and put*

$$r_\theta(D) = \frac{\theta}{1+\theta}\psi(D).$$

For every $0 < \zeta < 1$, there are constants $K_{\theta,\varepsilon,\zeta}, B, D_{\theta,\varepsilon} > 0$ such that every (C, D)-dense set S, with $0 < D \le D_{\theta,\varepsilon}$, has

$$\mathrm{Pull}(S) \text{ is } \left(K_{\theta,\varepsilon,\zeta}(C+1)D^{-B}, \frac{\zeta r_\theta(D)}{\log 2}\right)\text{-dense.} \tag{3.12}$$

After $\theta, \varepsilon, \zeta$ are fixed, the constants $K_{\theta,\varepsilon,\zeta}$, B, and $D_{\theta,\varepsilon}$ are independent of C, D, S, M, and of every later endpoint-bootstrap stage. Consequently every

$$0 < \chi < a_0 \frac{\theta}{1+\theta} \tag{3.13}$$

is an admissible exponent update $D \mapsto \chi D$ for all sufficiently small D.

The proof is given in Section 5 and is used in Theorem 1.1 through the endpoint-only recurrence of Section 6.

3.1. Parity-word and composition reduction

The first-odd-position decomposition used below is exact.

Lemma 3.10 (Parity-word/composition dictionary)**.** *Let $w \in \{0,1\}^M$ have weight $s \ge 1$, with odd positions*

$$0 \le i_1 < i_2 < \cdots < i_s < M.$$

Put

$$u = i_1, \qquad k_j = i_{j+1} - i_j \ (1 \le j < s), \qquad k_s = M - i_s. \tag{3.14}$$

Then $\mathbf{k} = (k_1, \dots, k_s)$ is a positive composition of $N = M - u$. Conversely, $u \ge 0$ and a positive composition of $M - u$ recover a unique weight-s word by

$$i_m = u + k_1 + \cdots + k_{m-1}. \tag{3.15}$$

Thus weight-s words with first odd position u are in bijection with positive compositions of $M - u$ into s parts.

If c_w is the affine correction from (2.6), then

$$c_w = \sum_{m=1}^{s} 3^{s-m} 2^{i_m} = 2^u \sum_{m=1}^{s} 3^{s-m} 2^{k_1+\cdots+k_{m-1}}. \tag{3.16}$$

Proof. The first equality follows by unwinding the correction recursion, and the second follows from (3.15). □

For a positive composition

$$\mathbf{k} = (k_1, \dots, k_s), \qquad k_1 + \cdots + k_s = N,$$

put $K_0 = 0$, $K_j = k_1 + \cdots + k_j$, and define the positive integer Syracuse numerator

$$\mathscr{A}_{\mathbf{k}} = \sum_{j=0}^{s-1} 3^{s-1-j} 2^{K_j}. \tag{3.17}$$

The normalized offset is

$$C_s(\mathbf{k}) = 2^{-N} \mathscr{A}_{\mathbf{k}} \pmod{3^s}. \tag{3.18}$$

Equation (3.16) says precisely that $c_w = 2^u \mathscr{A}_{\mathbf{k}}$. Since $M = u + N$, multiplication by 2^{-M} gives

$$2^{-M} c_w = 2^{-N} \mathscr{A}_{\mathbf{k}} = C_s(\mathbf{k}) \pmod{3^s}. \tag{3.19}$$

Let

$$L_{s,N} = \binom{N-1}{s-1}$$

and let $\nu_{s,N}$ be the uniform offset law of all positive compositions of N into s parts. Define its normalized Rényi-$(1+\theta)$ functional by

$$\mathcal{A}_{s,N}(\theta) = 3^{s\theta} \sum_{a \bmod 3^s} \nu_{s,N}(a)^{1+\theta}. \tag{3.20}$$

Multiplication by the unit $2^{-N} \bmod 3^s$ merely permutes residue classes. Thus the same functional is obtained from the fibers of $\mathscr{A}_{\mathbf{k}} \bmod 3^s$, as used in Section 3. Finally, put

$$\beta_{s,N} = L_{s,N} 2^{-N}$$

and

$$\mathfrak{A}_M(\theta) = 2^{-M} + \sum_{u=0}^{M-1} 2^{-u} \sum_{s=1}^{M-u} \beta_{s,M-u} \mathcal{A}_{s,M-u}(\theta). \tag{3.21}$$

We now prove the bridge to the actual Collatz endpoint fibers. Fix M, s, and split a weight-s parity word according to the position u of its first odd letter. Lemma 3.10 identifies the cohort with the positive compositions of $N = M - u$, so it has size $L_{s,N}$ and correction law $\nu_{s,N}$.

Let $f_{M,s,u}(y)$ be the endpoint multiplicity in this cohort, and let $c_{s,N}(a) = L_{s,N}\nu_{s,N}(a)$ be its correction multiplicity. If two starts in the same parity level have equal output, the affine identity

$$2^M T^M(n) = 3^s n + c_w$$

implies that their corrections are congruent modulo 3^s. Thus every endpoint fiber lies inside one correction class; by (3.19), these are exactly the normalized offset classes counted by $\nu_{s,N}$. More explicitly, for each correction class a, let $\mathcal{Y}_a$ be the endpoint fibers contained in that class. Then

$$\sum_{y \in \mathcal{Y}_a} f_{M,s,u}(y) \le c_{s,N}(a).$$

Since the summands are nonnegative and $r \ge 1$,

$$\sum_{y \in \mathcal{Y}_a} f_{M,s,u}(y)^r \le \left(\sum_{y \in \mathcal{Y}_a} f_{M,s,u}(y) \right)^r \le c_{s,N}(a)^r.$$

Summing over a gives

$$\sum_y f_{M,s,u}(y)^r \le \sum_{a \bmod 3^s} c_{s,N}(a)^r. \tag{3.22}$$

Write

$$f_{M,s}(y) = \sum_u f_{M,s,u}(y), \qquad \lambda_u = \frac{L_{s,M-u}}{\binom{M}{s}}.$$

The first-odd-position cohorts partition all weight-s words, so $\sum_u L_{s,M-u} = \binom{M}{s}$ and hence $\sum_u \lambda_u = 1$. Weighted convexity gives

$$f_{M,s}(y)^{1+\theta} \le \sum_u \lambda_u^{-\theta} f_{M,s,u}(y)^{1+\theta}.$$

Using (3.22) in each cohort and then (3.20),

$$\sum_y f_{M,s}(y)^{1+\theta} \le \binom{M}{s}^{\theta} 3^{-s\theta} \sum_u L_{s,M-u}\mathcal{A}_{s,M-u}(\theta). \tag{3.23}$$

Substitution into the definition of $G_M(\theta)$, with $\mathcal{Y}_s = 3^{s+1}$, gives

$$G_M(\theta) \le 3^{\theta}\mathfrak{A}_M(\theta). \tag{3.24}$$

Indeed, $2^{-M}L_{s,M-u} = 2^{-u}\beta_{s,M-u}$. The all-even word contributes $3^{\theta}2^{-M}$, and the coefficients in $\mathfrak{A}_M$ have total mass one because

$$\sum_{s=1}^{N} \beta_{s,N} = \frac{1}{2}, \qquad 2^{-M} + \frac{1}{2}\sum_{u=0}^{M-1} 2^{-u} = 1. \tag{3.25}$$

The first identity is the binomial theorem:

$$\sum_{s=1}^{N} \beta_{s,N} = 2^{-N}\sum_{s=1}^{N}\binom{N-1}{s-1} = 2^{-N}2^{N-1} = \frac{1}{2}.$$

The second then follows from $\sum_{u=0}^{M-1} 2^{-u} = 2(1-2^{-M})$.

It remains to prove the critical estimate

$$\mathfrak{A}_M\left(\frac{1}{2}\right) \le \frac{9}{4}M. \tag{3.26}$$

3.2. Integer-lift injectivity

Lemma 3.11 (Integer-lift injectivity)**.** *For fixed s, N, the map*

$$\mathbf{k} \longmapsto \mathscr{A}_{\mathbf{k}} \in \mathbb{N}$$

is injective.

Proof. Divide (3.17) by 2^N:

$$2^{-N}\mathscr{A}_{\mathbf{k}} = 3^{s-1}2^{-N} + 3^{s-2}2^{-(N-k_1)} + \cdots + 2^{-k_s}. \tag{3.27}$$

Because every coefficient 3^{s-1-j} is odd, the 2-adic valuation of the j-th summand is exactly $K_j - N$. Since $K_0 < K_1 < \cdots < K_{s-1}$, the first summand has the unique smallest 2-adic valuation, $-N$. After subtracting it, the unique smallest valuation is $-(N-k_1)$, recovering k_1. Repeating recovers $k_2, \ldots, k_s$. □

The use of rational 2-adic valuations in (3.27) is legitimate after clearing the common denominator 2^N; uniqueness of the minimum prevents cancellation at every stage.

3.3. *Quotient-order estimate*

For a mod 3^s, let f_a be the number of compositions with

$$\mathscr{A}_{\mathbf{k}} \equiv a \pmod{3^s}.$$

Lemma 3.12 (Quotient-order estimate)**.** *For every $\theta > 0$,*

$$\mathcal{A}_{s,N}(\theta) \le (1+\theta)\,\mathbb{E}_{s,N}\left[\left(\frac{3^s + \mathscr{A}_{\mathbf{k}}}{L_{s,N}}\right)^{\theta}\right], \tag{3.28}$$

where the expectation is uniform over the $L_{s,N}$ compositions.

Proof. Fix $a \in \{0, \dots, 3^s - 1\}$. By Lemma 3.11, the integer lifts in this fiber are distinct. Write them in increasing quotient order as

$$\mathscr{A}_i = a + 3^s q_i, \qquad 0 \le q_1 < \cdots < q_{f_a}.$$

Then $q_i \ge i - 1$, and

$$\begin{aligned}
\frac{f_a^{1+\theta}}{1+\theta} &= \int_0^{f_a} x^\theta\, dx \\
&\le \sum_{i=1}^{f_a} i^\theta \\
&\le \sum_{i=1}^{f_a} (1+q_i)^\theta \\
&\le 3^{-s\theta} \sum_{i=1}^{f_a} (3^s + \mathscr{A}_i)^\theta.
\end{aligned}$$

Multiply by $3^{s\theta}$, sum over a, and divide by $L_{s,N}^{1+\theta}$. □

4. Reverse-spend moments and fixed-total Rényi estimates

The injectivity and quotient-order results of the preceding section reduce endpoint overload to moments of the normalized correction Z_N. Reading a composition from its terminal end exposes the cumulative dyadic spend U_j: a larger terminal spend suppresses the corresponding term $3^{j-1}2^{-U_j}$, while prescribing a terminal block of total spend D has the exact geometric cost 2^{-D}. The resulting reproduction series is subcritical below its critical moment. After conditioning on a central number of parts, the same mechanism remains subcritical throughout a fixed central window and leaves only the polynomial factor in (5.25).

Choose each of the $N-1$ possible separators independently with probability $1/2$. This gives a uniform random positive composition of N. Its number of parts satisfies

$$J_N = 1 + \mathrm{Bin}(N-1, \tfrac{1}{2}), \qquad p_{N,s} := \Pr(J_N = s) = \frac{L_{s,N}}{2^{N-1}}. \tag{4.1}$$

Put

$$B_N(\theta) = \sum_{s=1}^{N} p_{N,s}\mathcal{A}_{s,N}(\theta) \tag{4.2}$$

and

$$Z_N = \frac{3^{J_N} + \mathscr{A}_{\mathbf{k}}}{2^N}. \tag{4.3}$$

Since $L_{s,N} = 2^{N-1} p_{N,s}$, Lemma 3.12 gives

$$B_N(\theta) \le (1+\theta) 2^\theta \mathbb{E}\left[p_{N,J_N}^{-\theta} Z_N^\theta \right]. \tag{4.4}$$

For $0 < \theta < 1/2$, concavity gives

$$\mathbb{E} p_{N,J_N}^{-2\theta} = \sum_{s=1}^{N} p_{N,s}^{1-2\theta} \le N^{2\theta}. \tag{4.5}$$

Thus

$$\left(\mathbb{E} p_{N,J_N}^{-2\theta}\right)^{1/2} \le N^\theta. \tag{4.6}$$

4.1. Reverse-spend moments

Reverse the composition and put

$$U_j = k_{J_N - j+1} + \cdots + k_{J_N}.$$

Reindexing (3.17) gives

$$Z_N = \sum_{j=1}^{J_N - 1} 3^{j-1} 2^{-U_j} + 4 \cdot 3^{J_N - 1} 2^{-N}. \tag{4.7}$$

Lemma 4.1 (Reverse-spend moments)**.** *For each fixed* $0 < r_0 < 1$*, there is* $K(r_0) < \infty$ *such that*

$$\sup_{N \ge 1} \mathbb{E} Z_N^{r_0} \le K(r_0). \tag{4.8}$$

Proof. For $0 < r \le 1$, subadditivity gives

$$(x_1 + \cdots + x_m)^r \le \sum_i x_i^r.$$

For $j \ge 1$, fix positive terminal parts $d_1, \ldots, d_j$, with total $D < N$. The probability that these are the final j parts and another part precedes them is exactly 2^{-D}. Among the final $D-1$ internal separator locations, the prescribed parts force $j-1$ separators and $D-j$ nonseparators. Requiring another part before this terminal block forces the immediately preceding separator. Thus exactly D independent Bernoulli separator bits are prescribed: j separators and $D-j$ nonseparators, with probability 2^{-D}. Therefore

$$\begin{aligned} \mathbb{E} \sum_{j=1}^{J_N-1} \left(3^{j-1} 2^{-U_j}\right)^r &\le 3^{-r} \sum_{j \ge 1} \left(3^r \sum_{d \ge 1} 2^{-(1+r)d} \right)^j \\ &= 3^{-r} \sum_{j \ge 1} m(r)^j, \end{aligned} \tag{4.9}$$

where

$$m(r) = \frac{3^r}{2^{1+r} - 1}. \tag{4.10}$$

Strict concavity of $x \mapsto x^r$ gives

$$\frac{1+3^r}{2} < 2^r \qquad (0 < r < 1),$$

so $m(r) < 1$, and (4.9) is bounded by

$$3^{-r}\frac{m(r)}{1-m(r)}. \tag{4.11}$$

For the terminal term, (4.1) gives

$$\mathbb{E}\left(4\cdot 3^{J_N-1}2^{-N}\right)^r = 2^r\left(\frac{1+3^r}{2^{1+r}}\right)^{N-1} \le 2^r. \tag{4.12}$$

Combining equations (4.7), (4.9), and (4.12) proves the lemma. □

The strict inequality $r_0 < 1$ is the only restriction in this lemma. In particular, for every fixed $0 < \theta < 1/2$, it applies with $r_0 = 2\theta$.

Lemma 4.2 (Critical first moments)**.** *For every $N \ge 1$,*

$$\mathbb{E}p_{N,J_N}^{-1} = N, \qquad \mathbb{E}Z_N = 2 + \frac{N-1}{4}. \tag{4.13}$$

Proof. Every $p_{N,s}$ is positive, so

$$\mathbb{E}p_{N,J_N}^{-1} = \sum_{s=1}^{N} p_{N,s}p_{N,s}^{-1} = N.$$

For the nonterminal part of (4.7), fixing j terminal parts of total $D < N$ has probability 2^{-D}, while the summand itself contributes $3^{j-1}2^{-D}$. There are $\binom{D-1}{j-1}$ positive compositions of D into j parts. Therefore

$$\begin{aligned}\mathbb{E}\sum_{j=1}^{J_N-1}3^{j-1}2^{-U_j} &= \sum_{D=1}^{N-1}4^{-D}\sum_{j=1}^{D}\binom{D-1}{j-1}3^{j-1}\\ &= \sum_{D=1}^{N-1}4^{-D}(1+3)^{D-1}\\ &= \frac{N-1}{4}.\end{aligned}$$

The terminal term has exact mean

$$4\,2^{-N}\mathbb{E}3^{J_N-1} = 4\,2^{-N}\left(\frac{1+3}{2}\right)^{N-1} = 2.$$

This proves (4.13). □

4.2. Fixed-total Rényi estimate

Fix $0 < \theta < 1/2$. Cauchy–Schwarz in (4.4), followed by (4.6) and Lemma 4.1 with $r_0 = 2\theta$, gives a constant $C_\theta^{(B)} < \infty$ such that

$$\begin{aligned} B_N(\theta) &\le (1+\theta)2^\theta \left(\mathbb{E}p_{N,J_N}^{-2\theta}\right)^{1/2} \left(\mathbb{E}Z_N^{2\theta}\right)^{1/2} \\ &\le C_\theta^{(B)} N^\theta. \end{aligned} \tag{4.14}$$

Because $\sum_s \beta_{s,N}\mathcal{A}_{s,N}(\theta) = B_N(\theta)/2$, (3.21) becomes

$$\mathfrak{A}_M(\theta) = 2^{-M} + \frac{1}{2}\sum_{u=0}^{M-1} 2^{-u} B_{M-u}(\theta). \tag{4.15}$$

The weights in (4.15) sum to one. Since $M-u \le M$, (4.14) gives

$$\mathfrak{A}_M(\theta) \le C_\theta^{(\mathfrak{A})}(M+1)^\theta, \qquad 0<\theta<1/2. \tag{4.16}$$

For an explicit admissible constant, put

$$K_Z(r) = 3^{-r}\frac{m(r)}{1-m(r)} + 2^r \qquad (0<r<1).$$

The proof above permits

$$C_\theta^{(B)} = (1+\theta)2^\theta K_Z(2\theta)^{1/2}, \qquad C_\theta^{(\mathfrak{A})} = 1 + C_\theta^{(B)},$$

and hence

$$C_\theta = 3^\theta C_\theta^{(\mathfrak{A})}$$

in Theorem 3.1.

At the critical endpoint, (4.4) and Lemma 4.2 give

$$\begin{aligned} B_N\left(\frac{1}{2}\right) &\le \frac{3}{2}\sqrt{2}\left(\mathbb{E}p_{N,J_N}^{-1}\right)^{1/2} (\mathbb{E}Z_N)^{1/2} \\ &\le \frac{3\sqrt{2}}{4}\sqrt{N(N+7)} \le \frac{9}{4}N \qquad (N\ge 2). \end{aligned} \tag{4.17}$$

The last inequality uses $N+7 \le (9/2)N$ for $N \ge 2$, with equality at $N=2$. For $N=1$, the unique composition gives $B_1(1/2) = \sqrt{3} \le 9/4$, so the final bound in (4.17) holds for every $N \ge 1$. Consequently,

$$\begin{aligned} \mathfrak{A}_M\left(\frac{1}{2}\right) &\le 2^{-M} + \frac{9}{8}\sum_{u=0}^{M-1} 2^{-u}(M-u) \\ &= \frac{9}{4}(M-1) + \frac{13}{4}2^{-M} \\ &\le \frac{9}{4}M. \end{aligned} \tag{4.18}$$

Here

$$\sum_{u=0}^{M-1} 2^{-u}(M-u) = 2M - 2 + 2^{1-M};$$

this follows by subtracting

$$\sum_{u=0}^{M-1} u2^{-u} = 2 - (M+1)2^{1-M}$$

from $M\sum_{u=0}^{M-1} 2^{-u} = 2M - M2^{1-M}$. The last inequality in (4.18) follows from $(13/4)2^{-M} \leq 13/8 \leq 9/4$. Together with (3.24), this proves the critical estimate

$$G_M\left(\frac{1}{2}\right) \leq \frac{9\sqrt{3}}{4} M. \tag{4.19}$$

For later use with a coarse polynomial recurrence, (4.19) immediately implies the weaker bound

$$G_M\left(\frac{1}{2}\right) \leq 9\sqrt{3}\, M\sqrt{M}, \tag{4.20}$$

because $M \geq 1$ and $9\sqrt{3}/4 \leq 9\sqrt{3}\sqrt{M}$. This is (3.3); the additional square-root factor changes no exponential rate below.

5. Nonlinear endpoint pullback

Proof of Theorem 3.4 and Corollaries 3.5–3.7. Put

$$\psi(D) = -\log\left(\frac{1+3^{-D}}{2}\right). \tag{5.1}$$

For $v > 0$, Markov's inequality and (4.19) show that the source mass carried by fibers with $g_{M,s}(y) > e^{vM}$ is at most

$$\frac{9\sqrt{3}}{4} Me^{-vM/2}. \tag{5.2}$$

Let S be (C, D)-dense, $0 < D \leq 1$. At parity level s, every output lies below 3^{s+1}, so the number of bad possible outputs is at most $C3^{(s+1)(1-D)}$. On fibers satisfying $g_{M,s}(y) \leq e^{vM}$, the normalized bad source mass is therefore at most

$$Ce^{vM}2^{-M}\sum_{s=0}^{M}\binom{M}{s}3^{-D(s+1)}.$$

The level sum is exact:

$$\begin{aligned} 2^{-M}\sum_{s=0}^{M}\binom{M}{s}3^{-D(s+1)} &= 3^{-D}\left(\frac{1+3^{-D}}{2}\right)^{M} \\ &= 3^{-D}e^{-\psi(D)M}. \end{aligned} \tag{5.3}$$

Thus the ordinary-fiber contribution is

$$C3^{-D}e^{-(\psi(D)-v)M}. \tag{5.4}$$

Choosing

$$v = \frac{2\psi(D)}{3} \tag{5.5}$$

balances (5.2) and (5.4). In fact it maximizes $\min(v/2, \psi(D) - v)$, so the factor $1/3$ is optimal within this two-term truncation. Hence an absolute K satisfies

$$\frac{\#\{n \in [2^M, 2^{M+1}) : T^M(n) \notin S\}}{2^M} \leq K(C+1)(M+1)e^{-\psi(D)M/3}. \tag{5.6}$$

The weaker endpoint bound (4.20) similarly gives

$$\frac{\#\{n \in [2^M, 2^{M+1}) : T^M(n) \notin S\}}{2^M} \leq K(C+1)(M+1)^2 e^{-\psi(D)M/3}. \tag{5.7}$$

Fix $0 < \eta < 1$. Since $0 < \psi(D) < 1$, putting $u = \psi(D)M$ gives

$$(M+1)e^{-(1-\eta)\psi(D)M/3} \leq \psi(D)^{-1}(u+1)e^{-(1-\eta)u/3}.$$

The last factor is bounded uniformly for $u \geq 0$. Thus there is K_η such that

$$(M+1)e^{-\psi(D)M/3} \leq K_\eta \psi(D)^{-1} e^{-\eta\psi(D)M/3}. \tag{5.8}$$

The shell-to-global lemma applies because $\eta\psi(D)/3 < \log 2$. Its geometric-series denominator is uniformly bounded away from zero on $0 < D \leq 1$. We obtain

$$\mathrm{Pull}(S) \text{ is } \left(K_\eta(C+1)\psi(D)^{-1}, \frac{\eta\psi(D)}{3\log 2}\right)\text{-dense.} \tag{5.9}$$

Using the weaker shell bound (5.7) instead proves the coarse estimate

$$\mathrm{Pull}(S) \text{ is } \left(K'_\eta(C+1)\psi(D)^{-2}, \frac{\eta\psi(D)}{3\log 2}\right)\text{-dense.} \tag{5.10}$$

Direct differentiation gives

$$\psi'(D) = \frac{\log 3}{3^D + 1}, \qquad \psi''(D) = -\frac{(\log 3)^2\, 3^D}{(3^D+1)^2} < 0. \tag{5.11}$$

Therefore ψ is concave, $\psi(0) = 0$, and $\psi(1) = \log(3/2)$. For $0 < D \leq 1$,

$$\psi(D) \geq D \log \frac{3}{2}. \tag{5.12}$$

Substituting this into (5.9) gives the sharper D^{-1} paper prefactor. Substitution into (5.10) proves Corollary 3.5 and gives Corollary 3.6: every

$$0 < \xi < c_{\mathrm{unif}} := \frac{\log(3/2)}{3\log 2} = 0.194987500240\ldots \tag{5.13}$$

is a valid linear density-update constant after changing the prefactor.

Finally, (5.11) gives

$$\lim_{D\downarrow 0} \frac{\psi(D)}{D} = \psi'(0) = \frac{\log 3}{2}. \tag{5.14}$$

Hence every

$$0 < \xi < c_{\mathrm{asym}} := \frac{\log 3}{6\log 2} = 0.264160416786\ldots \tag{5.15}$$

is valid in (5.9), and hence also in (5.10), for all sufficiently small D, after fixing $\eta < 1$ sufficiently close to one. This threshold may be made explicit. Given $0 < c_{\mathrm{L}} < c_{\mathrm{asym}}$, choose

$$\frac{c_{\mathrm{L}}}{c_{\mathrm{asym}}} < \eta < 1$$

and put

$$D_c = \min\left\{1, \frac{\log\left(\dfrac{\eta \log 3}{3c_{\mathrm{L}} \log 2} - 1\right)}{\log 3}\right\} > 0.$$

Since ψ' is decreasing, for $0 < D \le D_c$,

$$\frac{\eta\psi(D)}{3\log 2} \ge \frac{\eta D\psi'(D)}{3\log 2} = \frac{\eta D \log 3}{3\log 2(3^D + 1)} \ge c_{\mathrm{L}} D.$$

Neither endpoint (5.13) nor (5.15) is attained. This proves Theorem 3.4 and Corollaries 3.5–3.7. □

Proof of Theorems 3.8 and 3.9. We prove the cohort-wise estimate in a central window fixed independently of the shell size and the density exponent. Fix

$$\frac{1}{2} < \theta < 1, \qquad \alpha_\theta = \frac{2^\theta - 1}{3^\theta - 1}, \qquad 0 < \varepsilon < \alpha_\theta - \frac{1}{2}. \tag{5.16}$$

Two elementary inequalities locate the available central window:

$$\frac{1}{2} < \alpha_\theta < \frac{\log 2}{\log 3}. \tag{5.17}$$

For the first, strict concavity of x^θ gives $2^\theta > (1 + 3^\theta)/2$. For the second, the function

$$x \longmapsto \frac{x^\theta - 1}{\log x}$$

is strictly increasing on $(1, \infty)$: after writing $z = \theta \log x$, the derivative has the sign of $e^z(z-1) + 1$, whose derivative is $ze^z > 0$. These observations give (5.17).

Fix $N \ge 2$, $1 \le s \le N$, and put

$$\pi = \frac{s-1}{N-1}.$$

Under the uniform law on positive compositions of N into s parts, prescribe final parts $d_1, \dots, d_r$, of total $D = d_1 + \cdots + d_r < N$, and require another part before them. The exact probability is

$$\frac{\binom{N-1-D}{s-1-r}}{\binom{N-1}{s-1}}. \tag{5.18}$$

It also satisfies

$$\Pr_{s,N}(d_1, \dots, d_r \text{ terminal}) \le N\pi^r(1-\pi)^{D-r}. \tag{5.19}$$

Indeed, condition independent Bernoulli(π) separator bits on having exactly $s-1$ separators. The numerator of the resulting conditional probability is at most one. Put $n' = N - 1$ and $k' = s - 1 = n'\pi$. For

$0 < k' < n'$, consecutive binomial masses satisfy

$$\frac{\Pr(X = j + 1)}{\Pr(X = j)} = \frac{n' - j}{j + 1} \frac{k'}{n' - k'}.$$

At $j = k' - 1$ this ratio is $(n' - k' + 1)/(n' - k') > 1$, while at $j = k'$ it is $k'/(k' + 1) < 1$; the same formula shows monotonicity on either side. Hence $k' = s - 1$ is a mode of $\mathrm{Bin}(N - 1, \pi)$. The largest of its N masses is at least $1/N$. This proves (5.19); the endpoint cases $s = 1, N$ are deterministic and satisfy the same bound.

For a fixed composition reverse the parts and put

$$U_r = k_{s-r+1} + \cdots + k_s, \qquad Z_{s,N} = \frac{3^s + \mathscr{A}_{\mathbf{k}}}{2^N}.$$

The same reindexing as in (4.7) gives the exact identity

$$Z_{s,N} = \sum_{r=1}^{s-1} 3^{r-1} 2^{-U_r} + 4 \cdot 3^{s-1} 2^{-N}. \tag{5.20}$$

Let

$$\mathcal{I}_{N,\varepsilon} = \left\{ s : \left| s - \frac{N+1}{2} \right| \le \varepsilon N \right\}.$$

Choose π_* strictly between $1/2+\varepsilon$ and α_θ. For all sufficiently large N, every $s \in \mathcal{I}_{N,\varepsilon}$ has $\pi \le \pi_*$. Since

$$R_* = \frac{3^\theta \pi_*}{2^\theta - 1 + \pi_*} < 1, \tag{5.21}$$

subadditivity of x^θ, followed by (5.19), gives

$$\begin{aligned} \mathbb{E}_{s,N} \sum_{r=1}^{s-1} (3^{r-1} 2^{-U_r})^\theta &\le N 3^{-\theta} \sum_{r \ge 1} \left[3^\theta \pi \sum_{d \ge 1} 2^{-\theta d} (1 - \pi)^{d-1} \right]^r \\ &\le N 3^{-\theta} \frac{R_*}{1 - R_*}. \end{aligned} \tag{5.22}$$

The bracket evaluates to $3^\theta \pi / (2^\theta - 1 + \pi)$, which is increasing in π. By (5.17), there is a fixed $N_0 = N_0(\varepsilon)$ such that every $N \ge N_0$ and $s \in \mathcal{I}_{N,\varepsilon}$ satisfies $s/N \le \log_3 2$, so the final term in (5.20) is at most $4/3$. For the finite set

$$\mathcal{F}_\varepsilon = \{(s, N) : 1 \le N < N_0,\ s \in \mathcal{I}_{N,\varepsilon}\},$$

enlarge $K_{\varepsilon,\theta}$ by the maximum of $N^{-1} \mathbb{E}_{s,N} Z_{s,N}^\theta$ over $\mathcal{F}_\varepsilon$ (with the maximum interpreted as zero if the set is empty). Hence

$$\mathbb{E}_{s,N} Z_{s,N}^\theta \le K_{\varepsilon,\theta} N \qquad (s \in \mathcal{I}_{N,\varepsilon}). \tag{5.23}$$

Because $L_{s,N} = 2^{N-1} p_{N,s}$, the quotient-order estimate (3.28) and (5.23) imply

$$\mathcal{A}_{s,N}(\theta) \le C_{\varepsilon,\theta} N p_{N,s}^{-\theta}. \tag{5.24}$$

Multiplying by $p_{N,s}$, summing over the central window, and using concavity gives

$$\sum_{s \in \mathcal{I}_{N,\varepsilon}} p_{N,s} \mathcal{A}_{s,N}(\theta) \le C_{\varepsilon,\theta} N \sum_{s=1}^{N} p_{N,s}^{1-\theta} \le C_{\varepsilon,\theta} N^{1+\theta}. \tag{5.25}$$

Display (5.25) proves the fixed-total estimate required in Theorem 3.8.

Central endpoint bridge and discarded cohorts.

For fixed M, s, retain only first-odd-position cohorts u for which $s \in \mathcal{I}_{M-u,\varepsilon}$, and call their total endpoint multiplicity $f^{\mathrm{cen}}_{M,s}(y)$. The inherited weights of the retained cohorts sum to at most one. Adjoining a dummy zero cohort of the missing weight, or equivalently retaining the factor $(\sum_u \lambda_u)^\theta \le 1$, gives the same weighted-convexity bound as (3.23):

$$G^{\mathrm{cen}}_M(\theta,\varepsilon) \le 3^\theta \sum_{u=0}^{M-1} 2^{-u} \sum_{s\in\mathcal{I}_{M-u,\varepsilon}} \beta_{s,M-u}\mathcal{A}_{s,M-u}(\theta), \tag{5.26}$$

where

$$G^{\mathrm{cen}}_M(\theta,\varepsilon) = 2^{-M} \sum_{s,y} f^{\mathrm{cen}}_{M,s}(y) \left(\frac{3^{s+1} f^{\mathrm{cen}}_{M,s}(y)}{\binom{M}{s}}\right)^\theta .$$

Since $\beta_{s,N} = p_{N,s}/2$, (5.25) and the geometric sum over u yield

$$G^{\mathrm{cen}}_M(\theta,\varepsilon) \le C'_{\varepsilon,\theta}(M+1)^{1+\theta}. \tag{5.27}$$

The omitted source mass is exponentially small. At fixed u, its normalized mass is

$$2^{-u-1}\,\Pr\left(\left|J_{M-u} - \frac{M-u+1}{2}\right| > \varepsilon(M-u)\right).$$

For $N = M-u \ge 2$, the symmetric-binomial estimate from Section 2 gives

$$\Pr\left(\left|\mathrm{Bin}(N-1,1/2) - \frac{N-1}{2}\right| > \varepsilon N\right) \le 2e^{-2\varepsilon^2 N^2/(N-1)} \le 2e^{-2\varepsilon^2 N}.$$

Thus, after adding the all-even word, the discarded mass is at most

$$2^{-M} + \sum_{u=0}^{M-2} 2^{-u-1} 2e^{-2\varepsilon^2(M-u)}.$$

Choose $0 < c_\varepsilon < \min(2\varepsilon^2, \log 2)$. For every summand,

$$2^{-u}e^{-2\varepsilon^2(M-u)} \le e^{-c_\varepsilon M}e^{-(\log 2 - c_\varepsilon)u}.$$

The resulting geometric series is bounded independently of M, and $2^{-M} \le e^{-c_\varepsilon M}$. Consequently

$$\mathrm{mass}_{\mathrm{discard}}(M) \le C_\varepsilon e^{-c_\varepsilon M}. \tag{5.28}$$

Central nonlinear pullback.

Let S be (C,D)-dense. For $\nu > 0$, call a central endpoint fiber heavy when

$$\frac{3^{s+1} f^{\mathrm{cen}}_{M,s}(y)}{\binom{M}{s}} > e^{\nu M}.$$

By (5.27), the normalized source mass in heavy central fibers is at most

$$C_{\varepsilon,\theta}(M+1)^{1+\theta}e^{-\theta\nu M}. \tag{5.29}$$

On the remaining central fibers, the calculation (5.3)–(5.4) is unchanged and gives bad source mass at most

$$C3^{-D}e^{-(\psi(D)-\nu)M}. \tag{5.30}$$

The choice

$$\nu = \frac{\psi(D)}{1+\theta}$$

balances the two rates at

$$r_\theta(D) = \frac{\theta}{1+\theta}\psi(D). \tag{5.31}$$

Together with (5.28), this proves

$$\frac{\#\{n \in [2^M, 2^{M+1}) : T^M(n) \notin S\}}{2^M} \le C_\varepsilon e^{-c_\varepsilon M} + K_{\varepsilon,\theta}(C+1)(M+1)^{1+\theta}e^{-r_\theta(D)M}. \tag{5.32}$$

For sufficiently small D, one has $r_\theta(D) < c_\varepsilon$. Fix $0 < \zeta < 1$. We use the following elementary uniform absorption inequality: for $B \ge 0$, $M \ge 1$, and $u > 0$,

$$(M+1)^B e^{-uM} \le K_{B,\zeta}u^{-B}e^{-\zeta uM}, \qquad K_{B,\zeta} = \begin{cases} 1, & B = 0, \\ \left(\dfrac{2B}{e(1-\zeta)}\right)^B, & B > 0. \end{cases}$$

Indeed, $u(M+1) \le 2uM$, and the maximum of $x^B e^{-(1-\zeta)x}$ over $x > 0$ occurs at $x = B/(1-\zeta)$ when $B > 0$. Applying this with $u = r_\theta(D)$ gives

$$(M+1)^{1+\theta}e^{-(1-\zeta)r_\theta(D)M} \ll_{\theta,\zeta} r_\theta(D)^{-(1+\theta)}.$$

Since $r_\theta(D) \asymp_\theta D$, the fixed-rate dyadic summation lemma turns (5.32) into the following global estimate. Indeed, after decreasing the fixed cutoff in D if necessary, the discard term is bounded by $C_\varepsilon e^{-r_\theta(D)M}$, while the polynomial absorption gives

$$\frac{\#\{n \in [2^M, 2^{M+1}) : T^M(n) \notin S\}}{2^M} \le K_{\varepsilon,\theta,\zeta}(C+1)r_\theta(D)^{-(1+\theta)}e^{-\zeta r_\theta(D)M}.$$

We also shrink the cutoff so that $\zeta r_\theta(D) < \log 2$. The fixed-rate summation formula from Section 2 then contributes the finite factor

$$\frac{2}{2e^{-\zeta r_\theta(D)} - 1}.$$

This factor is uniformly bounded after one more fixed reduction of the small-D cutoff: for example, $\zeta r_\theta(D) \le \log(4/3)$ makes it at most 4. The finitely many shells before the shell estimate becomes valid are absorbed into the same density constant. Since $r_\theta(D) \asymp_\theta D$, we obtain

$$\mathrm{Pull}(S) \text{ is } \left(K_{\varepsilon,\theta,\zeta}(C+1)D^{-(1+\theta)}, \frac{\zeta r_\theta(D)}{\log 2}\right)\text{-dense}. \tag{5.33}$$

Because $0 < D \le 1$ and $1+\theta < 2$, this implies (3.12). Also, (5.14) gives

$$\lim_{D\downarrow 0}\frac{r_\theta(D)}{D\log 2} = a_0\frac{\theta}{1+\theta}. \tag{5.34}$$

The cutoff implicit in the final consequence can be written explicitly. Write $\Lambda(\theta) = a_0\theta/(1+\theta)$. Fix $0 < \chi < \Lambda(\theta)$, and then choose $0 < \zeta < 1$ with $\chi < \zeta\Lambda(\theta)$. Put

$$D_{\mathrm{slope}}(\theta, \chi, \zeta) = \min\left\{1, \frac{\log\left(\dfrac{\zeta\theta\log 3}{(1+\theta)\chi\log 2} - 1\right)}{\log 3}\right\} > 0. \tag{5.35}$$

The strict inequality $\chi < \zeta\Lambda(\theta)$ makes the quantity inside the logarithm greater than one. Since ψ' is decreasing and $\psi(0) = 0$, one has $\psi(D) \geq D\psi'(D)$. Therefore, for $0 < D \leq D_{\mathrm{slope}}(\theta, \chi, \zeta)$,

$$\begin{aligned} \frac{\zeta r_\theta(D)}{\log 2} &\geq \frac{\zeta\theta D\psi'(D)}{(1+\theta)\log 2} \\ &= \frac{\zeta\theta D\log 3}{(1+\theta)\log 2(3^D+1)} \geq \chi D. \end{aligned} \tag{5.36}$$

Thus every strict χ in (3.13) is available for

$$0 < D \leq \min\left\{D_{\theta,\varepsilon}, D_{\mathrm{slope}}(\theta, \chi, \zeta)\right\},$$

where $D_{\theta,\varepsilon}$ is the fixed cutoff already obtained in (5.33). This proves the final consequence of Theorem 3.9 with an explicit cutoff for the additional linear-slope comparison. Neither $\theta = 1$ nor the limiting transport constant is asserted. □

6. Endpoint bootstrap and proof of the main theorem

An orbit-minimum statement does not require concatenating a two-sided envelope through every intermediate iterate. For $0 < t \leq 1$, let W_t be the set of integers n satisfying

$$\rho^k n^{1-t} \leq T^k(n) \leq \rho^k n^{1+t} \qquad (0 \leq k \leq \lfloor \log_2 n \rfloor).$$

Proposition A.3 proves that there are absolute constants $K_W, c_W > 0$ for which W_t is $(K_W, c_W t^2)$-dense. We now iterate only logarithmic-block endpoints. Define

$$F(n) = T^{\lfloor \log_2 n \rfloor}(n). \tag{6.1}$$

6.1. A fixed contraction set

Fix

$$a_0 < r < 1, \qquad t = r - a_0. \tag{6.2}$$

Use the initial-window set W_t above and put

$$K_0 = 2^{1-a_0}.$$

If $M = \lfloor \log_2 n \rfloor$, then

$$\rho^M = 2^{-(1-a_0)M}, \qquad n < 2^{M+1},$$

and hence

$$\rho^M n^{1-a_0} < K_0.$$

The upper envelope at $k = M$ therefore gives

$$F(n) \le K_0 n^r \qquad (n \in W_t). \tag{6.3}$$

Define

$$\mathcal{E}_r = \{n : F(n) \le K_0 n^r\}. \tag{6.4}$$

Since $W_t \subseteq \mathcal{E}_r$, Proposition A.3 makes both sets (C_t, D_t)-dense for fixed positive constants depending only on r. For the iteration we retain the stronger set W_t, because its all-prefix envelope also controls the orbit between consecutive endpoints. No tolerance in this construction changes with the number of blocks.

6.2. Recursive endpoint sets

Put

$$\mathcal{S}_1 = W_t, \qquad \mathcal{S}_{j+1} = W_t \cap \mathrm{Pull}(\mathcal{S}_j). \tag{6.5}$$

If $n \in \mathcal{S}_R$, define $n_0 = n$ and $n_{i+1} = F(n_i)$. Induction in (6.5) shows that $n_i \in W_t \subseteq \mathcal{E}_r$ for $0 \le i < R$. A second induction using (6.3) gives

$$n_i \le K_0^{1+r+\cdots+r^{i-1}} n^{r^i} \le K_0^{1/(1-r)} n^{r^i}. \tag{6.6}$$

Each n_i is an actual Collatz iterate of n, so

$$T_{\min}(n) \le n_R. \tag{6.7}$$

More precisely, put

$$\tau_R = \sum_{i=0}^{R-1} \lfloor \log_2 n_i \rfloor. \tag{6.8}$$

Then $n_R = T^{\tau_R}(n)$. Taking logarithms in (6.6), summing over $0 \le i < R$, and using $\log K_0 = b \log 2$ gives

$$\tau_R \le \frac{\log n}{(1-r)\log 2} + \frac{bR}{1-r}. \tag{6.9}$$

This retains the actual iterate witnessing the endpoint contraction; the second term is only logarithmic in the outer shell under the schedule below.

The stronger choice of $\mathcal{S}_j$ also controls the entire orbit up to that witness. Let τ_i denote the accumulated time through the first i blocks. If $0 \le i < R$ and $0 \le \ell \le \lfloor \log_2 n_i \rfloor$, then membership $n_i \in W_t$, (6.6), and $\rho \le 1$ give

$$\begin{aligned} T^{\tau_i+\ell}(n) = T^{\ell}(n_i) &\le \rho^{\ell} n_i^{1+t} \\ &\le K_0^{(1+t)/(1-r)} n^{(1+t)r^i} \le K_0^{(1+t)/(1-r)} n^{1+t}. \end{aligned}$$

The same bound at τ_R follows from (6.6). Consequently

$$\max_{0 \le k \le \tau_R} T^k(n) \le K_0^{(1+t)/(1-r)} n^{1+t}. \tag{6.10}$$

Proposition 6.1 (Abstract endpoint-only bootstrap). *Fix $a_0 < r < 1$. Suppose that there are constants*

$$0 < \chi < a_0, \qquad B \ge 0, \qquad K_{\mathrm{L}}, D_c > 0, \tag{6.11}$$

such that, whenever $0 < D \le D_c$*, pullback sends every* (C, D)*-dense set to a set with density parameters*

$$(C, D) \longmapsto \big(K_{\mathrm{L}}(C+1)D^{-B}, \chi D\big). \tag{6.12}$$

Then every

$$0 < \delta < \frac{\log(1/r)}{\log(1/\chi)}$$

is admissible in the coefficient-one orbit-minimum conclusion (1.2)*. The same construction gives the witnessing-time estimate* (6.9) *and the quantitative shell estimate* (6.20)*. Because the recursive sets retain* W_t*, it also gives the orbit ceiling* (6.10)*.*

Proof. Choose $0 < D_1 \le \min(D_t, D_c, 1)$. After degrading the density of W_t, put $C_1 = C_t$. Since $0 < \chi < a_0 < 1$, induction gives $D_j \le D_1 \le D_c$, so the assumed pullback is applicable at every stage. Moreover $\chi D_j \le D_1 \le D_t$, so intersecting the pullback with W_t retains exponent χD_j. Consequently the sets in (6.5) are (C_j, D_j)-dense with

$$D_{j+1} = \chi D_j, \qquad C_{j+1} \le K_{\mathrm{L}}(C_j+1)D_j^{-B} + C_t. \tag{6.13}$$

In particular,

$$D_j = D_1\chi^{j-1}. \tag{6.14}$$

Specifically, take

$$A = \max\{1, K_{\mathrm{L}} + C_t + 2\}.$$

Since $D_j \le 1$, (6.13) then gives

$$C_{j+1} + 2 \le A(C_j+2)D_j^{-B}.$$

Taking logarithms and using (6.14),

$$\begin{aligned}\log(C_R+2) \le \log(C_1+2) + (R-1)\big(\log A + B\log D_1^{-1}\big)\\ + \frac{B}{2}(R-1)(R-2)\log(\chi^{-1}).\end{aligned}$$

Thus

$$\log(C_R+2) = O(R^2). \tag{6.15}$$

This is the only prefactor accumulation in the endpoint iteration.

Shell-dependent iteration.
For the shell $[2^M, 2^{M+1})$, put

$$R_M = \lceil \omega\log(M+4)\rceil \tag{6.16}$$

with fixed $\omega > 0$. Equation (6.14) gives

$$D_{R_M} \ge D_1(M+4)^{-\omega\log(1/\chi)}. \tag{6.17}$$

The density estimate for $\mathcal{S}_{R_M}$ bounds its exceptional proportion on shell M by

$$2C_{R_M}e^{-(\log 2)D_{R_M}M}. \tag{6.18}$$

By (6.15)–(6.17), this tends to zero provided

$$\omega\log(1/\chi) < 1. \tag{6.19}$$

Indeed, with $\gamma = \omega \log(1/\chi) < 1$, equations (6.15)–(6.17) give

$$\log(2C_{R_M}) - (\log 2) D_{R_M} M \leq O((\log M)^2) - cM^{1-\gamma} \longrightarrow -\infty$$

for some fixed $c > 0$. This is the logarithm of the right-hand side of (6.18), and hence proves its convergence to zero. After decreasing c and increasing the fixed starting shell, the same calculation gives the quantitative form

$$2C_{R_M} e^{-(\log 2) D_{R_M} M} \leq 2 \exp\left(-c(M+4)^{1-\gamma}\right), \qquad \gamma = \omega \log(1/\chi) < 1. \tag{6.20}$$

The varying-rate dyadic summation lemma then shows that

$$\bigcup_{M \geq 0} \left(S_{R_M} \cap [2^M, 2^{M+1})\right) \tag{6.21}$$

has natural density one.

On this set, (6.6) and (6.16) give

$$\log n_{R_M} \leq \frac{\log K_0}{1-r} + (M+1) \log 2 \, (M+4)^{-\omega \log(1/r)}. \tag{6.22}$$

Consequently

$$n_{R_M} \leq \exp\left((\log n)^{1-\delta}\right) \tag{6.23}$$

for all sufficiently large M whenever

$$\delta < \omega \log(1/r). \tag{6.24}$$

To see the coefficient one explicitly, divide the second term on the right of (6.22) by $(\log n)^{1-\delta}$. Since $M \asymp \log n$, the quotient is

$$O\left(M^{\delta - \omega \log(1/r)}\right) = o(1)$$

under (6.24). The fixed term $\log K_0/(1-r)$ is also $o((\log n)^{1-\delta})$. Here $\delta < 1$ follows from $\chi < a_0 < r$, (6.19), and (6.24). Thus the whole right-hand side of (6.22) is at most $(\log n)^{1-\delta}$ after increasing the fixed starting shell, which is exactly (6.23). Conditions (6.19) and (6.24) admit ω exactly when

$$\delta < \frac{\log(1/r)}{\log(1/\chi)}. \tag{6.25}$$

This proves Proposition 6.1. □

6.3. Critical-moment specialization

Choose $0 < \chi < a_0/3$. Theorem 3.4 and Corollary 3.7 supply the hypothesis (6.12) of Proposition 6.1 with $B = 2$ for all sufficiently small D.

Letting $r \downarrow a_0$ and $\chi \uparrow a_0/3$ in (6.25) gives every

$$\delta < \frac{\log(1/a_0)}{\log(3/a_0)} = 0.174719543136035\ldots. \tag{6.26}$$

Given any strict $\delta < \delta_{\mathrm{crit}}$, continuity permits fixed strict r, χ, and then an ω satisfying (6.19) and (6.24). Equations (6.7), (6.21), and (6.23) prove the critical-moment baseline. The boundary in (6.26) is a supremum and is not asserted at equality. The same recurrence, with the stronger transport supplied by Theorem 3.9, proves Theorem 1.1 in Section 6. The time accounting (6.8)–(6.9) and the quantitative shell estimate (6.20) apply unchanged after that replacement.

6.4. Moment phase transition

For $0 < \theta < 1$, define

$$\Lambda(\theta) = a_0 \frac{\theta}{1+\theta}, \qquad \alpha_*(\theta) = \frac{\log(1/a_0)}{\log(1/\Lambda(\theta))}. \tag{6.27}$$

The first quantity is the small-D transport slope obtained by balancing a $(1+\theta)$-moment tail against the ordinary-fiber term. The second is the endpoint-only bootstrap ceiling associated with that slope.

Lemma 6.2 (Phase transition)**.** *The functions Λ and α_* are strictly increasing on $(0, 1)$, and*

$$\Lambda\left(\frac{1}{2}\right) = \frac{a_0}{3}, \qquad \alpha_*\left(\frac{1}{2}\right) = \frac{\log(1/a_0)}{\log(3/a_0)} = 0.174719543136035\ldots. \tag{6.28}$$

Moreover,

$$\lim_{\theta\uparrow 1} \alpha_*(\theta) = \frac{\log(1/a_0)}{\log(2/a_0)} = 0.251245530155874\ldots. \tag{6.29}$$

Proof. The map $\theta \mapsto \theta/(1+\theta)$ is strictly increasing, so the same is true of Λ. For $0 < \theta < 1$,

$$0 < \Lambda(\theta) < \frac{a_0}{2} < 1.$$

Hence $\log(1/\Lambda(\theta))$ is positive and strictly decreasing. Since $\log(1/a_0) > 0$, the quotient defining α_* is strictly increasing. Substitution of $\theta = 1/2$ proves (6.28), and continuity as $\theta \uparrow 1$ proves (6.29). □

For $\theta > 0$ and $0 < \pi < 1$, define the reverse-spend reproduction factor and the unrestricted reverse-spend ratio by

$$\mathscr{R}_\theta(\pi) = \frac{3^\theta \pi}{2^\theta - 1 + \pi}, \qquad m(\theta) = \frac{3^\theta}{2^{1+\theta} - 1}, \tag{6.30}$$

and retain $\alpha_\theta = (2^\theta - 1)/(3^\theta - 1)$ from Section 5.

Proposition 6.3 (Reverse-spend critical curve)**.** *For every $\theta > 0$ and $0 < \pi < 1$,*

$$\mathscr{R}_\theta(\pi) < 1 \quad \Longleftrightarrow \quad \pi < \alpha_\theta. \tag{6.31}$$

At the centered proportion $\pi = 1/2$, the two quantities are identical:

$$\mathscr{R}_\theta(1/2) = \frac{3^\theta/2}{2^\theta - 1/2} = \frac{3^\theta}{2^{1+\theta} - 1} = m(\theta).$$

Their common value is less than one for $0 < \theta < 1$, equal to one at $\theta = 1$, and greater than one for $\theta > 1$. Thus $(\theta, \pi) = (1, 1/2)$ is the simultaneous critical point of the centered cohort and unrestricted reverse-spend series. In particular, the restriction in the unrestricted reverse-spend estimate and the collapse of the central-window margin are the same critical phenomenon.

Proof. The denominator in $\mathscr{R}_\theta(\pi)$ is positive, and direct cross-multiplication gives

$$\mathscr{R}_\theta(\pi) < 1 \quad \Longleftrightarrow \quad (3^\theta - 1)\pi < 2^\theta - 1,$$

which is (6.31). For $0 < \theta < 1$, strict concavity of x^θ at the midpoint $2 = (1+3)/2$ gives

$$2^\theta > \frac{1+3^\theta}{2}.$$

For $\theta > 1$, strict convexity reverses this inequality, while equality holds at $\theta = 1$. The comparison with $1/2$ in (6.31) follows immediately. The displayed identity with $m(\theta)$ then gives its trichotomy without a second argument. □

The central moment proof uses the fixed window

$$0 < \varepsilon < \varepsilon_0(\theta) = \frac{2^\theta - 1}{3^\theta - 1} - \frac{1}{2}. \tag{6.32}$$

Strict concavity of x^θ gives $2^\theta > (1+3^\theta)/2$, so $\varepsilon_0(\theta) > 0$ for every fixed $\theta < 1$. This interval shrinks to zero as $\theta \uparrow 1$, which is why the window must be chosen after fixing θ, but still independently of the density exponent D.

Thus $\varepsilon_0(\theta)$ is the exact half-width allowed by the uniform reproduction-factor inequality $\mathscr{R}_\theta(1/2 + \varepsilon) < 1$. It is not claimed to obstruct cohort-averaged, signed, or generated-target estimates using additional endpoint structure.

Theorem 6.4 (Central higher-moment endpoint)**.** *Fix $1/2 < \theta < 1$ and $0 < \varepsilon < \varepsilon_0(\theta)$. Then for every*

$$0 < \delta < \alpha_*(\theta) \tag{6.33}$$

one has

$$T_{\min}(n) \le \exp\left((\log n)^{1-\delta}\right) \tag{6.34}$$

on a set of natural density one.

Proof. Markov's inequality applied to Theorem 3.8 bounds the source mass in fibers with normalized overload greater than $e^{\nu M}$ by

$$C_{\theta,\varepsilon}(M+1)^{B_{\theta,\varepsilon}} e^{-\theta \nu M}.$$

The omitted cohorts have mass $O_\varepsilon(e^{-c_\varepsilon M})$ by the symmetric-binomial estimate proved in Section 2. On the remaining ordinary fibers, the calculation in (5.3)–(5.4) gives $C3^{-D}e^{-(\psi(D)-\nu)M}$. Choosing

$$\nu = \frac{\psi(D)}{1+\theta}$$

balances the two variable rates at $r_\theta(D) = \theta\psi(D)/(1+\theta)$. Absorbing the fixed polynomial factor into D^{-B}, and then applying the fixed-rate dyadic summation lemma, proves Theorem 3.9. Since

$$\lim_{D \downarrow 0} \frac{r_\theta(D)}{D \log 2} = \Lambda(\theta),$$

every strict $\chi < \Lambda(\theta)$ is available in the density recurrence (6.12). More explicitly, for a fixed strict $\delta < \alpha_*(\theta)$, choose fixed $r > a_0$ and $0 < \chi < \Lambda(\theta)$ such that $\delta < \log(1/r)/\log(1/\chi)$; then choose fixed $\zeta < 1$ with $\chi < \zeta\Lambda(\theta)$, choose the initial density exponent below all resulting fixed small-D cutoffs, and finally choose ω satisfying (6.19) and (6.24). None of these choices varies with the shell, target set, or bootstrap stage. Proposition 6.1 applies directly with the fixed prefactor power from Theorem 3.9. Letting $r \downarrow a_0$ and $\chi \uparrow \Lambda(\theta)$ in (6.25) gives exactly (6.33)–(6.34). □

To prove Theorem 1.1, fix $0 < \delta < \delta_{\text{end}}$. By (6.29) and strict inequality, choose one fixed $\theta < 1$ sufficiently close to one that $\delta < \alpha_*(\theta)$, and then choose one fixed $0 < \varepsilon < \varepsilon_0(\theta)$. Theorem 6.4 gives the required density-one conclusion. Thus (6.29) is an unconditional nonattained endpoint; neither $\theta = 1$ nor equality at δ_{end} is claimed.

For a concrete interior example, take

$$\delta = 0.24, \quad \theta = 0.95, \quad \varepsilon = 0.005, \quad \zeta = 0.999, \quad \chi = 0.385, \quad r = 0.7925, \quad \omega = 1.04. \tag{6.35}$$

Direct substitution gives

$$\varepsilon_0(0.95) = 0.00654820\ldots > 0.005, \qquad \zeta\Lambda(0.95) = 0.3856945\ldots > 0.385,$$

and

$$\omega \log(1/r) = 0.2418652\ldots > 0.24, \qquad \omega \log(1/\chi) = 0.9926924\ldots < 1.$$

Thus (6.19) and (6.24) hold. Moreover

$$\frac{1}{(1-r)\log 2} = 6.9527471\ldots < 6.953,$$

so this same example also satisfies the advertised witnessing-time margin.

7. Full-envelope concatenation and recurrence

For the full-envelope companion, retain the entire initial window W_t defined in Section 6. Proposition A.3 supplies its density; no quadratic pullback theorem is used here.

For $0 < \lambda < 1$ and $t > 0$, let $\mathcal{G}(\lambda, t)$ be the set of integers n, with $M = \lfloor \log_2 n \rfloor$, for which

$$\rho^k n^{1-t} \le T^k(n) \le \rho^k n^{1+t}$$

for every

$$0 \le k \le \left\lfloor \frac{(1-\lambda)M}{b} \right\rfloor.$$

Because $(1-a_0)/b = 1$, one has $\mathcal{G}(a_0, t) = W_t$, so Proposition A.3 supplies the initial case.

Lemma 7.1 (One-step envelope concatenation)**.** *Let* $0 < \lambda < 1$, $a_0 < q < 1$, $0 < t_{\mathrm{old}}, \xi \le 1$, *and* $t_{\mathrm{new}} > 0$. *Put*

$$\tau = t_{\mathrm{old}} + \xi + t_{\mathrm{old}}\xi, \qquad \mu = \lambda(q - a_0) - \xi(1-\lambda).$$

If $\tau \le t_{\mathrm{new}}$, $\mu > 0$, *and*

$$M = \lfloor \log_2 n \rfloor \ge \frac{1+b}{\mu},$$

then

$$n \in \mathrm{Pull}(\mathcal{G}(\lambda, t_{\mathrm{old}})) \cap W_\xi \quad \Longrightarrow \quad n \in \mathcal{G}(q\lambda, t_{\mathrm{new}}).$$

Proof. Put $m = T^M(n)$ and $M' = \lfloor \log_2 m \rfloor$. Membership in W_ξ gives

$$\rho^M n^{1-\xi} \le m \le \rho^M n^{1+\xi}.$$

For every i in the window supplied by $m \in \mathcal{G}(\lambda, t_{\mathrm{old}})$,

$$T^{M+i}(n) \le \rho^i (\rho^M n^{1+\xi})^{1+t_{\mathrm{old}}} \le \rho^{M+i} n^{1+\tau},$$

and similarly

$$T^{M+i}(n) \geq \rho^i(\rho^M n^{1-\xi})^{1-t_{\mathrm{old}}} \geq \rho^{M+i} n^{1-\tau}.$$

The first M iterates satisfy the same envelope because $\xi \leq \tau$.

Since $n \geq 2^M$,

$$m \geq 2^{(a_0-\xi)M}, \qquad M' \geq (a_0 - \xi)M - 1.$$

Consequently,

$$M + \left\lfloor \frac{(1-\lambda)M'}{b} \right\rfloor \geq \frac{(1-q\lambda)M}{b} + \frac{\mu M}{b} - \frac{1-\lambda}{b} - 1.$$

The lower bound on M makes the final three terms nonnegative, so the concatenated window reaches the horizon defining $\mathcal{G}(q\lambda, t_{\mathrm{new}})$. □

7.1. Parameterized linear envelope recurrence

We now plug Theorem 3.4 into Lemma 7.1. Fix

$$0 < \kappa < \sqrt{2} - 1, \qquad a_0 < q < 1. \tag{7.1}$$

Put

$$\lambda_j = a_0 q^j, \qquad t_j = \kappa t_{j+1}, \qquad c_q = \frac{q - a_0}{4\kappa q}. \tag{7.2}$$

As before, require

$$t_{j+1} \leq c_q \lambda_{j+1}. \tag{7.3}$$

The tolerance loss in Lemma 7.1 is

$$2t_j + t_j^2 \leq (2\kappa + \kappa^2)t_{j+1} < t_{j+1}, \tag{7.4}$$

provided $t_{j+1} \leq 1$. Also,

$$t_j \leq \frac{q - a_0}{4}\lambda_j, \tag{7.5}$$

so the positive-margin and finite-startup arguments are unchanged. The condition $\kappa q < 1$, used to propagate (7.3) backward, follows automatically from (7.1).

For a fixed finite schedule, write

$$\mathcal{G}_j = \mathcal{G}(\lambda_j, t_j).$$

Let $K_W, c_W > 0$ be the initial-window constants, and define

$$d_W = \min(1, c_W). \tag{7.6}$$

Fix any

$$0 < c_{\mathrm{L}} < c_{\mathrm{asym}}. \tag{7.7}$$

Choose η and the explicit $D_c > 0$ from the paragraph following (5.15). Theorem 3.4 then supplies $K_{\mathrm{L}} > 0$ such that the pullback exponent is at least $c_{\mathrm{L}} D$ and its prefactor is at most $K_{\mathrm{L}}(C+1)D^{-2}$ whenever $0 < D \leq D_c$. The sharper $\psi(D)^{-1} = O(D^{-1})$ prefactor of Theorem 3.4 would only improve the constants in (7.18); the bootstrap uses the conservative form in Corollary 3.5 to keep the recurrence uniform.

Suppose $\mathcal{G}_j$ is (C_j, D_j)-dense and

$$0 < D_j \le \min(D_c, d_W t_j^2). \tag{7.8}$$

By (7.5), the margin in Lemma 7.1 obeys

$$\mu_j \ge \frac{3}{4}(q - a_0)\lambda_j.$$

Thus it is enough to impose

$$M \ge M_{0,j} := \left\lceil \frac{4(1+b)}{3(q-a_0)\lambda_j} \right\rceil.$$

Lemma 7.1 therefore gives

$$\mathrm{Pull}(\mathcal{G}_j) \cap W_{t_j} \setminus [1, 2^{M_{0,j}}) \subseteq \mathcal{G}_{j+1}.$$

Equivalently,

$$\mathcal{G}_{j+1}^c \subseteq \left(\mathrm{Pull}(\mathcal{G}_j) \cap W_{t_j}\right)^c \cup [1, 2^{M_{0,j}}).$$

The finite interval contributes fewer than

$$2^{M_{0,j}} \le \exp(\Gamma_q/\lambda_j),$$

with Γ_q as in (7.11); here $\lambda_j < 1$ absorbs the ceiling. Corollary 3.5, Proposition A.3, and the elementary density calculus therefore give

$$D_{j+1} = c_{\mathrm{L}} D_j, \tag{7.9}$$

$$C_{j+1} \le K_{\mathrm{L}}(C_j + 1)D_j^{-2} + K_W + \exp(\Gamma_q/\lambda_j), \tag{7.10}$$

where

$$\Gamma_q = \left(\frac{4(1+b)}{3(q-a_0)} + 1\right)\log 2. \tag{7.11}$$

Since $c_{\mathrm{L}} < c_{\mathrm{asym}} < 1$, (7.8) gives

$$c_{\mathrm{L}} D_j \le D_j \le d_W t_j^2 \le c_W t_j^2,$$

and

$$D_{j+1} \le d_W t_j^2 \le d_W t_{j+1}^2,$$

so (7.8) is preserved.

At stage zero choose

$$C_0 = K_W, \qquad D_0 = d_W t_0^2. \tag{7.12}$$

For the shell-dependent schedule below, $t_0 \to 0$. Hence $D_0 \le D_c$ for all sufficiently large outer shells; no fixed positive lower bound on D is introduced. Then (7.9) is exact:

$$D_j = d_W t_0^2 c_{\mathrm{L}}^j. \tag{7.13}$$

Proposition 7.2 (Stage induction)**.** *For every finite schedule of length R satisfying (7.1)–(7.8), define C_j, D_j by (7.9)–(7.12). Then*

$$\mathcal{G}(\lambda_j, t_j) \text{ is } (C_j, D_j)\text{-dense} \qquad (0 \le j \le R).$$

Proof. At $j = 0$,

$$\mathcal{G}(\lambda_0, t_0) = \mathcal{G}(a_0, t_0) = W_{t_0},$$

because $(1 - a_0)/b = 1$. Proposition A.3 and $D_0 = d_W t_0^2 \le c_W t_0^2$ give the base case.

Assume the assertion at stage j. Corollary 3.5 makes $\mathrm{Pull}(\mathcal{G}_j)$

$$\big(K_{\mathrm{L}}(C_j + 1)D_j^{-2},\ c_{\mathrm{L}} D_j\big)\text{-dense.}$$

Proposition A.3 gives the density of W_{t_j}. Under (7.8), $c_{\mathrm{L}} D_j \le c_W t_j^2$, so the intersection lemma gives exponent $D_{j+1} = c_{\mathrm{L}} D_j$ and the first two terms of (7.10). The complement containment above and the finite-modification lemma supply the final term. This proves the inductive step. □

7.2. Prefactor accumulation

Put

$$L_j = \log(C_j + 2)$$

and

$$K_\dagger = \max\left(1,\ K_{\mathrm{L}} + \frac{K_W + 3}{2}\right). \tag{7.14}$$

Because $D_j \le 1$, $C_j + 2 \ge 2$, and $\exp(\Gamma_q/\lambda_j) \ge 1$, (7.10) implies

$$C_{j+1} + 2 \le K_\dagger (C_j + 2) D_j^{-2} \exp(\Gamma_q/\lambda_j).$$

Consequently

$$L_{j+1} \le L_j + \log K_\dagger + 2\log\frac{1}{D_j} + \frac{\Gamma_q}{\lambda_j}. \tag{7.15}$$

From (7.13),

$$\log\frac{1}{D_j} = \log\frac{1}{d_W} + 2\log\frac{1}{t_0} + j\log\frac{1}{c_{\mathrm{L}}}. \tag{7.16}$$

Summing (7.15), and using

$$\sum_{j=0}^{R-1}\frac{1}{\lambda_j} \le \frac{q}{(1-q)\lambda_R}, \tag{7.17}$$

gives the explicit bound

$$\begin{aligned} L_R \le L_0 + R\log K_\dagger + 2R\left(\log\frac{1}{d_W} + 2\log\frac{1}{t_0}\right) \\ + R(R-1)\log\frac{1}{c_{\mathrm{L}}} + \frac{\Gamma_q q}{(1-q)\lambda_R}. \end{aligned} \tag{7.18}$$

The conservative D_j^{-2} prefactor contributes only $2\sum_j \log(1/D_j) = O(R^2 + R\log(1/t_R))$. It does not change the exponent recurrence, and the finite-startup term remains dominant.

Since $t_0 = \kappa^R t_R$, equations (7.13) and (7.18) imply

$$D_R = d_W t_R^2 (c_{\mathrm{L}}\kappa^2)^R, \tag{7.19}$$

$$\log(C_R + 2) = O_{q,\kappa,c_{\mathrm{L}}}\left(\lambda_R^{-1} + R^2 + R\log\frac{1}{t_R}\right). \tag{7.20}$$

8. Power schedule, optimization, and quantitative consequences

Put

$$X_M = M + 4$$

and choose a fixed slope $\omega > 0$. Use the schedule

$$R_M = \lceil \omega \log X_M \rceil . \tag{8.1}$$

Define

$$\upsilon = \omega \log(1/q), \qquad w = \omega \log \left(\frac{1}{c_{\mathrm{L}} \kappa^2 q^2} \right). \tag{8.2}$$

The ceiling inequalities give

$$q X_M^{-\upsilon} \le q^{R_M} \le X_M^{-\upsilon} \tag{8.3}$$

and

$$(c_{\mathrm{L}} \kappa^2 q^2) X_M^{-w} \le (c_{\mathrm{L}} \kappa^2 q^2)^{R_M} \le X_M^{-w}. \tag{8.4}$$

For this outer shell, terminate the schedule at

$$\lambda_{R_M} = a_0 q^{R_M}, \qquad t_{R_M} = \frac{c_q \lambda_{R_M}}{2(1 + c_q \lambda_{R_M})}, \qquad t_j = \kappa^{R_M - j} t_{R_M}. \tag{8.5}$$

Apply the stage-induction proposition to this shell-dependent schedule, and denote its density constants by $C_j^{(M)}, D_j^{(M)}$. Within this section we suppress the superscript (M) to keep the estimates readable. Then

$$t_{R_M} \asymp_{q,\kappa} X_M^{-\upsilon}. \tag{8.6}$$

Indeed, (8.3) gives

$$a_0 q X_M^{-\upsilon} \le \lambda_{R_M} \le a_0 X_M^{-\upsilon},$$

and $0 < \lambda_{R_M} \le a_0$ gives

$$\frac{c_q}{2(1 + c_q a_0)} \lambda_{R_M} \le t_{R_M} \le \frac{c_q}{2} \lambda_{R_M}.$$

These two explicit sandwiches prove (8.6). The earlier stage condition (7.3) follows from

$$\frac{t_{j+1}}{\lambda_{j+1}} = \frac{t_{R_M}}{\lambda_{R_M}} (\kappa q)^{R_M - j - 1} \le c_q. \tag{8.7}$$

Equations (7.19) and (8.3)–(8.6) give a constant $d_-(q, \kappa, c_{\mathrm{L}}) > 0$ such that

$$D_{R_M} \ge d_-(q, \kappa, c_{\mathrm{L}}) X_M^{-w}. \tag{8.8}$$

For completeness, put $h = c_{\mathrm{L}} \kappa^2 q^2$, so $0 < h < 1$. From (8.5),

$$t_{R_M} \ge \frac{c_q a_0}{2(1 + c_q a_0)} q^{R_M},$$

while the ceiling bound gives $h^{R_M} \ge h X_M^{-w}$. Therefore one may take

$$d_-(q, \kappa, c_{\mathrm{L}}) = d_W \left(\frac{c_q a_0}{2(1 + c_q a_0)} \right)^2 h,$$

which proves (8.8) directly from (7.19). The explicit recurrence (7.18), together with (8.3)–(8.6), gives

$$\log(C_{R_M}+2) = O_{q,\kappa,c_{\mathrm{L}}}\left((\log X_M)^2 + X_M^{\upsilon}\log X_M\right). \tag{8.9}$$

Assume

$$\upsilon + w < 1. \tag{8.10}$$

Then the right-hand side of (8.9) is $o(X_M^{1-w})$, while $D_{R_M}M \gg X_M^{1-w}$. Consequently, for all sufficiently large M,

$$2C_{R_M}\exp(-D_{R_M}M\log 2) \le 2\exp(-c_{q,\kappa,c_{\mathrm{L}}}X_M^{1-w}). \tag{8.11}$$

Indeed, (8.8) and $M = X_M - 4$ imply, for all sufficiently large M,

$$D_{R_M}M \ge \frac{d_-}{2}X_M^{1-w}.$$

On the other hand, (8.9) and $\upsilon + w < 1$ give

$$\log(2C_{R_M}) = o(X_M^{1-w}).$$

After increasing the fixed starting shell, the latter is at most $(d_- \log 2/4)X_M^{1-w}$. Taking logarithms then gives

$$\log\left(2C_{R_M}e^{-D_{R_M}M\log 2}\right) \le -\frac{d_-\log 2}{4}X_M^{1-w},$$

which is (8.11) with $c_{q,\kappa,c_{\mathrm{L}}} = d_- \log 2/4$. Indeed, define the terminal good part of the M-th shell by

$$\mathscr{H}_M = \mathcal{G}(\lambda_{R_M}, t_{R_M}) \cap [2^M, 2^{M+1}).$$

The stage-induction proposition and the definition of (C,D)-density give

$$\frac{\#([2^M,2^{M+1})\setminus\mathscr{H}_M)}{2^M} \le 2C_{R_M}e^{-D_{R_M}M\log 2},$$

which is bounded by the right-hand side of (8.11). Choose M_* beyond the finitely many asymptotic validity thresholds above and put

$$\mathscr{H} = [1, 2^{M_*}) \cup \bigcup_{M\ge M_*}\mathscr{H}_M.$$

This is the shell-indexed good set used in Theorem 1.4; its complement is not obtained from one infinite recursion. Put $\mathscr{E} = \mathbb{N}_{\ge 1}\setminus\mathscr{H}$. Lemma 2.1 gives $\mathscr{H}$ natural density one and, with the same dyadic summation, the quantitative bound

$$\#(\mathscr{E}\cap[1,X]) \le C_{q,\kappa,c_{\mathrm{L}}}X\exp\left(-c_{q,\kappa,c_{\mathrm{L}}}(\log X)^{1-w}\right). \tag{8.12}$$

Now fix a requested exponent $0 < \delta < 1$ and put

$$\varepsilon_M = X_M^{-\delta}. \tag{8.13}$$

If

$$\delta < \upsilon, \tag{8.14}$$

then the terminal tolerance eventually dominates the floor and endpoint losses:

$$M(\lambda_{R_M} + t_{R_M}) + b + 1 + t_{R_M} \le M\varepsilon_M. \tag{8.15}$$

To make the asymptotic comparison explicit, (8.3), (8.5), and (8.6) give

$$M(\lambda_{R_M} + t_{R_M}) + b + 1 + t_{R_M} = O(MX_M^{-\upsilon} + 1).$$

After division by $M\varepsilon_M = MX_M^{-\delta}$, this is $O(X_M^{\delta-\upsilon} + X_M^{\delta-1}) = o(1)$, because $\delta < \upsilon$ and $\upsilon < 1$ by (8.10). This proves (8.15) beyond a fixed starting shell. For $n \in \mathscr{H}_M$, put

$$k^* = \left\lfloor \frac{(1-\lambda_{R_M})M}{b} \right\rfloor.$$

Since $\log_2 \rho = -b$ and $\log_2 n < M + 1$, the terminal upper envelope gives

$$\begin{aligned} \log_2 T^{k^*}(n) &\le -bk^* + (1 + t_{R_M}) \log_2 n \\ &\le M(\lambda_{R_M} + t_{R_M}) + b + 1 + t_{R_M} \\ &\le M\varepsilon_M \le \varepsilon_M \log_2 n. \end{aligned}$$

Therefore

$$T_{\min}(n) \le n^{\varepsilon_M}. \tag{8.16}$$

There is no base-conversion loss here:

$$2^{\varepsilon_M \log_2 n} = n^{\varepsilon_M}.$$

Since

$$\log n < (M+1)\log 2 \le M + 4 = X_M,$$

monotonicity of negative real powers gives

$$n^{X_M^{-\delta}} \le \exp\left((\log n)^{1-\delta}\right). \tag{8.17}$$

The two strict inequalities $\delta < \upsilon$ and $\upsilon + w < 1$ admit a slope $\omega > 0$ exactly when

$$\delta < \frac{\log(1/q)}{\log\left(1/(c_{\mathrm{L}}\kappa^2 q^3)\right)}. \tag{8.18}$$

Theorem 8.1 (Parameterized stretched-logarithmic descent)**.** *Fix*

$$a_0 < q < 1, \qquad 0 < \kappa < \sqrt{2} - 1, \qquad 0 < c_{\mathrm{L}} < c_{\mathrm{asym}}.$$

Then every

$$0 < \delta < \frac{\log(1/q)}{\log(1/(c_{\mathrm{L}}\kappa^2 q^3))} \tag{8.19}$$

is admissible in the conclusion of Theorem 1.4. More precisely, the good set has the witnessing-time bound of Corollary 1.3, and its exceptional count satisfies (8.12) *for any schedule slope ω chosen so that*

$$\frac{\delta}{\log(1/q)} < \omega < \frac{1}{\log(1/(c_{\mathrm{L}}\kappa^2 q^3))}. \tag{8.20}$$

Proof. The interval in (8.20) is nonempty exactly under (8.19). Its left inequality is (8.14), and its right inequality is (8.10). Equations (8.11)–(8.17) then give the density, exceptional-count, orbit, and witnessing-time conclusions. □

For example, the interior choice

$$q = 0.793, \qquad \kappa = 0.414, \qquad c_{\mathrm{L}} = 0.264$$

gives the right-hand side of (8.19) as $0.0611735\ldots$. Thus $\delta = 0.06$ is admissible, and one may take $\omega = 0.261$, which lies between the two endpoints $0.258696\ldots$ and $0.263756\ldots$ in (8.20).

8.1. Parameter frontier and theorem consequences

Lemma 8.2 (Parameter frontier)**.** *One has*

$$\sup_{\substack{q>a_0 \\ 0<\kappa<\sqrt{2}-1 \\ 0<c_{\mathrm{L}}<c_{\mathrm{asym}}}} \frac{\log(1/q)}{\log(1/(c_{\mathrm{L}}\kappa^2 q^3))} = \delta_{\mathrm{env}}. \tag{8.21}$$

Proof. Put

$$A = \log(1/q), \qquad H = \log\frac{1}{c_{\mathrm{L}}\kappa^2}.$$

The ratio in (8.21) is

$$\frac{A}{H+3A}. \tag{8.22}$$

It is strictly increasing in A, c_{L}, and κ. Since A decreases as q increases, the supremum is approached as

$$q \downarrow a_0, \qquad \kappa \uparrow \sqrt{2}-1, \qquad c_{\mathrm{L}} \uparrow c_{\mathrm{asym}}.$$

Substitution gives

$$\delta_{\mathrm{env}} = \left[3 + \frac{\log\left(1/(c_{\mathrm{asym}}(\sqrt{2}-1)^2)\right)}{\log(1/a_0)}\right]^{-1}. \tag{8.23}$$

□

The witnessing-time coefficient used below satisfies

$$\frac{1}{b\log 2} = \frac{2}{\log(4/3)} < \frac{6953}{1000}.$$

Indeed, the first two positive terms of $\log((1+x)/(1-x)) = 2\sum_{j\ge 0} x^{2j+1}/(2j+1)$ at $x = 1/7$ give $\log(4/3) > 296/1029$, and $1{,}029{,}000 < 148 \cdot 6953 = 1{,}029{,}044$.

Proof of Theorem 1.4. By Lemma 8.2, any $0 < \delta < \delta_{\mathrm{env}}$ permits fixed choices

$$q > a_0, \qquad \kappa < \sqrt{2}-1, \qquad c_{\mathrm{L}} < c_{\mathrm{asym}}$$

for which (8.19) holds. The stage-zero exponent tends to zero with the outer shell, so the small-D cutoff chosen after (7.7) is respected at every stage for all sufficiently large shells. Theorem 8.1 proves

$$T_{\min}(n) \le \exp\left((\log n)^{1-\delta}\right) \tag{8.24}$$

on a set of natural density one. It also supplies the shell failure bound (8.12). The witnessing-time bound follows from the explicit k^* calculation preceding (8.16) and the estimate just established.

For $n \in \mathscr{H}_M$, take $\lambda_n = \lambda_{R_M}$ and $t_n = t_{R_M}$. Membership in $\mathscr{H}_M = \mathcal{G}(\lambda_{R_M}, t_{R_M})$ is exactly the displayed two-sided envelope. Equations (8.3), (8.5), and (8.6) give $\lambda_n, t_n \ll (\log n)^{-\upsilon}$, and the density assertion follows from the construction of $\mathscr{H}$. □

Numerically,

$$\delta_{\mathrm{env}}^{-1} = 16.3023489963\ldots, \qquad \delta_{\mathrm{env}} = 0.06134085341\ldots. \tag{8.25}$$

If one uses only the uniform chord bound (5.12), the corresponding rigorous supremum is

$$\delta_{\mathrm{unif}} = 0.05679313823\ldots. \tag{8.26}$$

The stronger value in (8.25) uses the explicitly proved fact that the shell-dependent stage-zero exponent tends to zero.

The bound (8.24) is asymptotically stronger than

$$\exp\left(\frac{C_a \log n}{(\log\log n)^a}\right)$$

for every fixed $a > 0$.

Proof of Corollary 1.2. Fix $0 < \delta < \delta_{\mathrm{end}}$ and

$$0 < \sigma < 1 - \frac{\delta}{\delta_{\mathrm{end}}}.$$

Equivalently,

$$\frac{\delta}{1-\sigma} < \delta_{\mathrm{end}}.$$

By (6.29), and then by taking strict interior values $r > a_0$ and $\chi < \Lambda(\theta)$, choose fixed $1/2 < \theta < 1$, r, and χ such that

$$\frac{\delta}{1-\sigma} < \frac{\log(1/r)}{\log(1/\chi)}.$$

Put

$$\omega = \frac{1-\sigma}{\log(1/\chi)}.$$

Then

$$\delta < \omega \log(1/r), \qquad \omega \log(1/\chi) = 1 - \sigma < 1.$$

Thus the hypotheses (6.19) and (6.24) hold, and (6.20) bounds the shell exceptional proportion by

$$2\exp\left(-c_0(M+4)^{\sigma}\right)$$

for some $c_0 > 0$ and all sufficiently large M. Apply Lemma 2.1 with

$$A = 2, \qquad \gamma = c_0, \qquad \text{the chosen exponent } \sigma.$$

Its explicit choice $A' = 2A + 1$ gives the prefactor 5. Thus, after putting

$$c_{\delta,\sigma} = \min\left\{\frac{c_0}{(4\log 2)^{\sigma}}, \frac{1}{2}\right\}$$

and increasing $X_{\delta,\sigma}$ to absorb the finitely many startup shells, one obtains

$$\#\left\{1 \le n \le X : T_{\min}(n) > \exp\left((\log n)^{1-\delta}\right)\right\} \le 5X \exp\left(-c_{\delta,\sigma}(\log X)^{\sigma}\right) \tag{8.27}$$

for every $X \ge X_{\delta,\sigma}$.

Finally, the interval

$$\delta < \sigma < 1 - \frac{\delta}{\delta_{\text{end}}}$$

is nonempty exactly when $\delta < \delta_{\text{end}}/(1 + \delta_{\text{end}})$, proving the last assertion. □

Proof of Corollary 1.3. First note that

$$\frac{1}{b \log 2} = \frac{2}{\log(4/3)}.$$

Putting $x = 1/7$ in

$$\log \frac{1+x}{1-x} = 2 \sum_{j \ge 0} \frac{x^{2j+1}}{2j+1}$$

gives the certified lower bound

$$\log \frac{4}{3} > 2\left(\frac{1}{7} + \frac{1}{3 \cdot 7^3}\right) = \frac{296}{1029}.$$

Consequently,

$$\frac{2}{\log(4/3)} < \frac{1029}{148} < \frac{6953}{1000}. \tag{8.28}$$

The final strict rational inequality is 1,029,000 < 1,029,044.

Now fix $0 < \delta < \delta_{\text{end}}$, and, for the orbit-ceiling assertion, fix $\beta > 0$. By (6.29), choose a fixed $\theta < 1$ and then $0 < \chi < \Lambda(\theta)$ so that

$$\delta < \frac{\log(1/a_0)}{\log(1/\chi)}.$$

Take $r > a_0$ sufficiently close to a_0 that

$$\delta < \frac{\log(1/r)}{\log(1/\chi)} \quad \text{and} \quad \frac{1}{(1-r)\log 2} < 6.953, \qquad r - a_0 < \beta$$

all hold. This is possible by continuity and (8.28). Choose ω strictly between the two bounds in (6.19) and (6.24). For $R = R_M = O(\log M)$, the actual iterate time τ_R in (6.8) satisfies by (6.9)

$$\tau_R \le \frac{\log n}{(1-r)\log 2} + O_r(\log\log n) < 6.953 \log n$$

for all sufficiently large n. Equations (6.8) and (6.23) show that this iterate has the required value. The density-one conclusion is the same varying-shell assembly as in Theorem 6.4. Finally, (6.10) bounds the orbit through this witness by $K_* n^{1+r-a_0}$, where K_* is fixed. Since $r - a_0 < \beta$, the fixed factor is at most $n^{\beta-(r-a_0)}$ for all sufficiently large n, proving the asserted $n^{1+\beta}$ ceiling. □

Proof of Corollary 1.5. Put

$$c = -\log \rho = b \log 2, \qquad c_* = c^{-1} = \frac{2}{\log(4/3)}.$$

We first record the common first-passage calculation. Suppose Theorem 1.4 supplies $t_n, \lambda_n \to 0$, and let $0 < y_n < 1$ be a target exponent. Write $L = \log n$ and

$$H_n = \left\lfloor \frac{(1-\lambda_n)\lfloor \log_2 n \rfloor}{b} \right\rfloor.$$

Since $\lfloor \log_2 n \rfloor = L/\log 2 + O(1)$,

$$H_n = \frac{(1-\lambda_n)L}{c} + O(1).$$

If $(y_n - t_n - \lambda_n)L \to +\infty$, then

$$k_n^+ = \left\lceil \frac{(1+t_n-y_n)L}{c} \right\rceil$$

lies below H_n for all sufficiently large n. The upper envelope at k_n^+ gives $T^{k_n^+}(n) \le n^{y_n}$, so the corresponding first passage $\tau_{n^{y_n}}(n)$ is defined and is at most k_n^+. Applying the lower envelope at that first-passage time gives

$$\frac{(1-t_n-y_n)L}{c} \le \tau_{n^{y_n}}(n) \le \frac{(1+t_n-y_n)L}{c} + 1. \tag{8.29}$$

Choose the parameters in Theorem 1.4 and Section 8, and put $\sigma = 1-w$. Equations (8.10) and (8.12) give

$$\delta < \upsilon < \sigma < 1$$

and constants $C_0, c_0 > 0$ such that the complement of the envelope set $A_\delta = \mathscr{H}$ has cardinality at most $C_0 X \exp(-c_0(\log X)^\sigma)$ up to X, after increasing C_0 to absorb the fixed initial interval.

On A_δ, Theorem 1.4 gives $t_n, \lambda_n = O(L^{-\upsilon})$. Hence

$$\inf_{L^{-\delta} \le y \le 1} (y - t_n - \lambda_n)L \longrightarrow +\infty.$$

The argument leading to (8.29) therefore applies uniformly to every $L^{-\delta} \le y < 1$; for $y = 1$, the first-passage time is zero and the same conclusion is immediate. Since $c_* = c^{-1}$, equation (8.29) gives

$$\sup_{L^{-\delta} \le y \le 1} |\tau_{n^y}(n) - c_*(1-y)L| \le c_* t_n L + 1 = O(L^{1-\upsilon}),$$

which is (1.7).

Taking $y = L^{-\delta}$ and using $n^{L^{-\delta}} = \exp(L^{1-\delta})$ gives

$$\tau_{h_\delta}(n) = c_* L - c_* L^{1-\delta} + O(L^{1-\upsilon}).$$

For every fixed $0 < x < 1$, one has $L^{-\delta} \le x$ eventually, so the choice $y = x$ in (1.7) gives (1.8). The two ratio limits follow by division by L. □

Proof of Corollary 1.6. Fix $\beta > \delta_{\text{end}}^{-1}$. By (6.29), the central transport limit, and continuity, choose once and for all fixed parameters

$$\frac{1}{2} < \theta < 1, \qquad a_0 < r < 1, \qquad 0 < \chi < \Lambda(\theta)$$

such that

$$p := \frac{\log(1/\chi)}{\log(1/r)} < \beta. \tag{8.30}$$

Theorem 3.9 and Proposition 6.1 then provide a fixed initial density exponent $0 < D_1 \le 1$, fixed pullback constants, and sets $\mathcal{S}_R$ which are (C_R, D_R)-dense with

$$D_R = D_1 \chi^{R-1} \qquad (R \ge 1). \tag{8.31}$$

All these choices depend on β, but not on the power α.

Now fix $0 < \alpha < 1$ and put

$$R_\alpha = \left\lceil \frac{\log(1/\alpha)}{\log(1/r)} \right\rceil. \tag{8.32}$$

Then $R_\alpha \ge 1$, $r^{R_\alpha} \le \alpha$, and, because $R_\alpha - 1 < \log(1/\alpha)/\log(1/r)$,

$$D_{R_\alpha} = D_1 \chi^{R_\alpha - 1} \ge D_1 \alpha^p \ge D_1 \alpha^\beta. \tag{8.33}$$

For every $n \in \mathcal{S}_{R_\alpha}$, equation (6.6) gives

$$T_{\min}(n) \le K_0^{1/(1-r)} n^{r^{R_\alpha}} \le K_0^{1/(1-r)} n^\alpha.$$

Take

$$K_\beta = K_0^{1/(1-r)}, \qquad c_\beta = D_1, \qquad C_{\alpha,\beta} = C_{R_\alpha}.$$

The defining complement bound for the $(C_{R_\alpha}, D_{R_\alpha})$-dense set, followed by (8.33), is exactly (1.9). $\square$

9. Scope and literature comparison

For every fixed $\varepsilon > 0$, the orbit threshold in Theorem 1.1 is eventually smaller than n^ε. Natural density one implies logarithmic density one by partial summation: if $B(X) = o(X)$, then

$$\sum_{\substack{n \le X \\ n \text{ exceptional}}} \frac{1}{n} = \frac{B(X)}{X} + \int_1^X \frac{B(x)}{x^2}\, dx = o(\log X).$$

Thus Theorem 1.1 also holds in logarithmic density. Tao's result [7] permits an arbitrary function tending to infinity in that density notion and therefore has a much stronger orbit threshold. The present theorem instead supplies the stated stretched-logarithmic threshold in natural density. The bounded witnessing time and every exceptional-count exponent $\sigma < 1 - \delta/\delta_{\text{end}}$ hold throughout Theorem 1.1 range; in particular $\sigma > \delta$ is available for $\delta < 0.2007963458\ldots$. Theorem 1.4 retains the stronger full-segment conclusion in its smaller range $\delta < \delta_{\text{env}}$.

The main comparison axes are recorded in Table 1. The table is not a total ordering: the rates depend on different variables, and the displayed clocks count either half-map or raw-Collatz steps as indicated.

Table 1. *Selected comparison axes for almost-all Collatz descent.*

Work	Density and terminal threshold	Quantitative information	Structural feature
Allouche [2]	natural; n^θ, $\theta > 0.869\ldots$	no rate or clock compared here	fixed-power density estimate
Korec [5]	natural; n^θ, $\theta > 0.792\ldots$	no rate or clock compared here	sharper fixed-power threshold
Inselmann [4] (preprint)	natural; n^ε for each fixed $\varepsilon > 0$	horizon $2\log n/\log(4/3)$ in half-map steps	fixed-tolerance two-sided envelope
Tao [7]	logarithmic; arbitrary $f(n) \to \infty$	global theorem stated without a natural-density rate	first-passage stabilization and 3-adic mixing
Allikvere [1] (preprint)	natural; arbitrary $f(n) \to \infty$	fixed-target ratio $O((\log N_0)^{-1/29} + X^{-1/2000})$; raw clock $< 12\log n$	conditioned extension of Tao's framework
Mazur [6] (preprint)	natural; arbitrary $f(n) \to \infty$	odd fixed-target ratio $C_d(\log N_0)^{-d}$; raw clock $< 436\log n$	Tao-based transport with a Rhin phase input
This paper	natural; $\exp((\log n)^{1-\delta})$, $\delta < \delta_{\mathrm{end}}$	$5X\exp(-c(\log X)^\sigma)$; $< 6.953\log n$ half-map, hence $< 13.906\log n$ raw	independent fixed-total endpoint transport; shrinking envelope and first-passage profile in a smaller range

The fixed-target rates in [1, 6] and the counting-cutoff rate (1.3) are not estimates of the same quantity. The arbitrary-diverging-threshold conclusions in those preprints are more general on the terminal-threshold axis. The contributions isolated here are the independent fixed-total mechanism, the prescribed stretched-logarithmic rate with its counting-cutoff estimate, and the shrinking-tolerance pathwise conclusions.

Theorem 1.1 answers Question 1.5 of [4] positively for the explicit family $f_\delta(n) = \exp((\log n)^{1-\delta})$, with the quantitative rate (1.3). The remark following that question in [4] treats fixed powers and some unspecified functions that are $O(n^\theta)$ for every fixed $\theta > 0$, without stating this stretched-logarithmic family or an exceptional-count rate for it. A density-one diagonalization over fixed powers can produce some unspecified $\varepsilon(n) \to 0$, but its decay is governed by untracked onset thresholds; it does not supply the prescribed rate $\varepsilon(n) = (\log n)^{-\delta}$ or the exceptional count (1.3). The additional quantitative conclusions established here are the explicit threshold-dependent stretched-exponential exceptional count, the shrinking-error two-sided envelope, the uniform quantitative first-passage profile, and a self-contained fixed-total proof; these are separate comparison axes from arbitrary-diverging-threshold natural-density claims, not a global ordering among concurrent public manuscripts.

Remark 9.1 (Worst-window limitation)**.** Proposition 6.3 also identifies a precise limitation of a natural critical window argument. At $\theta = 1$,

$$\mathscr{R}_1\left(\frac{1}{2}+\varepsilon\right) = \frac{3(1/2+\varepsilon)}{3/2+\varepsilon}, \qquad \log \mathscr{R}_1\left(\frac{1}{2}+\varepsilon\right) = \frac{4}{3}\varepsilon + O(\varepsilon^2).$$

A proof that bounds every cohort in a window by its worst upper edge and discards the exterior only with a Hoeffding estimate must take $\varepsilon^2 \gtrsim D$ in order to make the discarded mass $e^{-\Omega(DM)}$. Its worst-edge reproduction loss is then $e^{\Omega(\sqrt{D}M)}$, which cannot be absorbed by a linear $e^{-\Theta(DM)}$ transport gain as $D \downarrow 0$. Thus this particular worst-window/Hoeffding continuation cannot attain the critical $\theta = 1$ slope. This is not an impossibility theorem for cohort-averaged, signed, or generated-target transport estimates.

The orbit-minimum conclusion does not exclude a hypothetical nontrivial cycle, nor does it show that almost every orbit avoids eventual entry into one. If C is any fixed cycle with minimum m_C, every point in its basin already satisfies $T_{\min}(n) \le m_C$; hence every sufficiently large point of that basin satisfies Theorem 1.1 automatically. A density-zero theorem for eventual entry would require an all-depth summable first-entry pullback estimate, which is not proved here. Pointwise cycle exclusion is instead the separate affine divisibility problem $(2^k - 3^s)n = c_w$, together with positivity and the cycle-minimum prefix conditions.

A. Maximal-barrier initial window

This appendix proves the initial-window proposition used in Section 6 without importing the quadratic pullback. For a Boolean parity word, put

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

Lemma A.1 (Two-sided maximal Boolean-walk bound)**.** *For $M \ge 1$ and $h \ge 0$,*

$$2^{-M} \# \left\{ w \in \{0,1\}^M : H_M(w) > h \right\} \le 2 \exp\left(-\frac{2h^2}{M} \right). \tag{A.2}$$

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

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

The identity

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

shows that the total potential is preserved when an unpruned node is expanded. The frontier sum is therefore $(\cosh\theta)^M$. Each crossing node contributes at least $2^{-|u|}\cosh(2\theta h)$, and its cylinder has measure $2^{-|u|}$. Hence

$$\Pr(H_M > h) \le \frac{(\cosh\theta)^M}{\cosh(2\theta h)} \le 2\exp\left(\frac{M\theta^2}{2} - 2\theta h \right).$$

Here we used $\cosh x \le e^{x^2/2}$ and $\cosh x \ge e^{|x|}/2$. Choosing $\theta = 2h/M$ proves (A.2); the case $h = 0$ is immediate. □

For the real affine correction, write

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

With

$$d_k = s_k - \frac{k}{2}, \qquad \rho = \frac{\sqrt{3}}{2},$$

the exact affine formula (2.7) is equivalently

$$T^k(n) = \rho^k n 3^{d_k} + \left(r_k(n) 3^{k/2} \right) 3^{d_k}. \tag{A.4}$$

Lemma A.2 (Uniform affine correction)**.** *If $H_M \le h$, then, for every $0 < k \le M$,*

$$r_k(n) 3^{k/2} \le (2+\sqrt{3}) \left(1 - \left(\frac{\sqrt{3}}{2} \right)^k \right) 3^h < (2+\sqrt{3}) 3^h. \tag{A.5}$$

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

$$\begin{aligned} r_k 3^{k/2} &\le 3^h \sum_{i=0}^{k-1} \frac{3^{(k-i-1)/2}}{2^{k-i}} \\ &= \frac{3^h}{2} \sum_{j=0}^{k-1} \left(\frac{\sqrt{3}}{2}\right)^j, \end{aligned}$$

which is (A.5). □

Proposition A.3 (Initial-window density)**.** *There are absolute constants $K_W, c_W > 0$ such that W_t is $(K_W, c_W t^2)$-dense for every $0 < t \le 1$.*

Proof. Put $L_3 = \log_2 3$ and choose

$$h = \frac{tM}{2L_3}.$$

Lemma A.1 gives

$$\Pr(H_M > h) \le 2\exp\left(-\frac{t^2 M}{2L_3^2}\right). \tag{A.6}$$

Assume $H_M \le h$, $M \ge 4$, and $tM \ge 2$. Since $|d_k| \le h$, one has

$$3^h = 2^{tM/2}, \qquad n^t \ge 2^{tM}, \qquad 2^{tM/2} \le \frac{1}{2} n^t. \tag{A.7}$$

The second relation uses $n \ge 2^M$, and the last uses it together with $tM \ge 2$. In particular,

$$3^{-h} = 2^{-tM/2} \ge n^{-t}, \qquad 3^h \le \frac{1}{2} n^t.$$

Therefore the multiplicative part of (A.4) satisfies

$$\rho^k n^{1-t} \le \rho^k n 3^{d_k} \le \frac{1}{2} \rho^k n^{1+t}. \tag{A.8}$$

For the additive part, Lemma A.2 gives

$$\left(r_k 3^{k/2}\right) 3^{d_k} \le (2+\sqrt{3}) 3^{2h} = (2+\sqrt{3}) 2^{tM}.$$

Because $k \le M$, $n \ge 2^M$, and

$$2^{4a_0} = 9 > 2(2+\sqrt{3}),$$

the condition $M \ge 4$ implies

$$(2+\sqrt{3}) 2^{tM} \le \frac{1}{2} \rho^k n^{1+t}. \tag{A.9}$$

Equations (A.8)–(A.9) prove the two-sided envelope throughout the first M iterates.

If $M < 4$ or $tM < 2$, use the trivial shell bound 1; in either case $t^2 M < 4$. Thus absolute $K_{\mathrm{sh}}, c_0 > 0$ satisfy

$$\frac{\#(W_t^c \cap [2^M, 2^{M+1}))}{2^M} \le K_{\mathrm{sh}} e^{-c_0 t^2 M}$$

for every shell. Applying the fixed-rate dyadic summation formula at the start of Section 2 uniformly over $0 < t \le 1$ proves Proposition A.3. Explicitly, choose

$$0 < c_0 < \min\left(\frac{1}{2L_3^2}, \log 2\right)$$

and $K_{\mathrm{sh}} = \max(2, e^{4c_0})$. The maximal-barrier estimate handles the large-shell case, while $1 \le e^{4c_0}e^{-c_0 t^2 M}$ handles $t^2 M < 4$. The dyadic summation denominator satisfies

$$2e^{-c_0 t^2} - 1 \ge 2e^{-c_0} - 1 > 0,$$

so

$$K_W = \frac{2K_{\mathrm{sh}}}{2e^{-c_0} - 1}, \qquad c_W = \frac{c_0}{\log 2}$$

works simultaneously for every $0 < t \le 1$. $\square$

Acknowledgments. The author thanks the contributors to Lean and Mathlib for the formal tools used in the supplementary verification.

Use of generative AI. The author used generative AI tools, principally OpenAI ChatGPT and Codex, for exploratory proof discussion, Lean assistance, literature-search assistance, and editorial review. Model outputs were treated as suggestions and not as mathematical evidence. The author independently verified every retained argument, computation, citation, and formal declaration and assumes full responsibility for the manuscript.

Funding statement. The author received no specific funding for this work.

Competing interests. The author declares no competing interests.

Data availability and reproducibility. No numerical experiment or external data set is used as a proof input. A supplementary Lean 4 development accompanies the paper at the frozen tag https://github.com/shaikidris/CET/tree/v2.0.1. A software release is identified by DOI https://doi.org/10.5281/zenodo.21797535. It is pinned to Lean `v4.15.0` and Mathlib commit `9837ca9d65d9de6fad1ef4381750ca688774e608`; running `lake build` from the `lean` directory builds the paper-facing chain. The declaration map and the separate dependency and axiom replay commands are recorded in `lean/CollatzEndpointTransport/README.md`. The axiom audit reports only `propext`, `Classical.choice`, and `Quot.sound`. The formalization is supplementary and is not used to discharge a step of the paper proof. The command `lake build` rebuilds every retained `Common` and `Linear` module; the narrower command `lake build CollatzEndpointTransport.Linear.PaperAudit` builds the public paper axiom-audit target.

References

[1] Jaan Allikvere. Almost all Collatz orbits attain almost bounded values in natural density, July 2026. URL https://doi.org/10.5281/zenodo.21499244. Preprint, version 2, 22 July 2026.

[2] Jean-Paul Allouche. Sur la conjecture de Syracuse–Kakutani–Collatz. In *Séminaire de Théorie des Nombres, 1978–1979*, pages 1–15. CNRS, Talence, 1979. Exp. no. 9.

[3] C. J. Everett. Iteration of the number-theoretic function $f(2n) = n$, $f(2n+1) = 3n+2$. *Advances in Mathematics*, 25(1): 42–45, 1977. http://dx.doi.org/10.1016/0001-8708(77)90087-1. URL https://doi.org/10.1016/0001-8708(77)90087-1.

[4] Manuel Inselmann. An approximation of the Collatz map and a lower bound for the average total stopping time, 2024. URL https://arxiv.org/abs/2402.03276. Version 3.

[5] Ivan Korec. A density estimate for the $3x+1$ problem. *Mathematica Slovaca*, 44(1):85–89, 1994. URL http://dml.cz/dmlcz/133225.

[6] Lech Mazur. Natural-density almost-bounded Collatz orbits in logarithmic time, July 2026. URL https://www.proofatlas.ai/papers/natural-density-log-time-collatz/Mazur_Natural_Density_Collatz_Orbits_in_Logarithmic_Time_v2.pdf. Preprint, version 2, 21 July 2026.

[7] Terence Tao. Almost all orbits of the Collatz map attain almost bounded values. *Forum of Mathematics, Pi*, 10:e12, 2022. http://dx.doi.org/10.1017/fmp.2022.8. URL https://doi.org/10.1017/fmp.2022.8.

[8] Riho Terras. A stopping time problem on the positive integers. *Acta Arithmetica*, 30(3):241–252, 1976. http://dx.doi.org/10.4064/aa-30-3-241-252. URL https://doi.org/10.4064/aa-30-3-241-252.