Natural-Density Collatz Descent to a Stretched-Logarithmic Scale

Idris Ali Shaik
Independent researcher
shaikidris@apache.org
ORCID: 0009-0009-9699-9712

Abstract

Let

$$T(n) = \begin{cases} n/2, & n \text{ even}, \\ (3n+1)/2, & n \text{ odd}, \end{cases} \qquad T_{\min}(n) = \min_{k \geq 0} T^k(n).$$

Put $a_0 = (\log_2 3)/2$ and

$$\delta_0 = \frac{\log(1/a_0)}{\log(2/a_0)} = 0.251245530155874\ldots.$$

For every $0 < \delta < \delta_0$, we prove that

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

for a set of positive integers of natural density one. We also give a stretched-exponential bound for the exceptional count and locate a descent witness before $6.953 \log n$ iterations of T. The proof conditions a dyadic parity word on its total number of odd steps, applies a Rényi estimate to the resulting composition fibers when the number of parts is near its mean, and iterates only the endpoints of logarithmic-length blocks. The result is an almost-all statement and does not prove the Collatz conjecture.

Keywords: Collatz map, natural density, Rényi moment, parity word, endpoint fiber, quantitative descent.
MSC 2020: Primary 11B83; Secondary 60F10.

1. Introduction

The Collatz conjecture asks whether every positive integer eventually reaches 1 under the map T above. This paper addresses a different and weaker question: how small a value can be proved to occur for almost every initial integer, when "almost every" means ordinary natural density?

Natural-density descent separates two questions that the full conjecture combines: how deeply a typical orbit can be forced to descend, and whether every orbit terminates. Replacing a known threshold by a smaller one measures deeper certified descent while retaining ordinary counting density. It does not, by itself, say anything about every individual orbit.

Throughout, log denotes the natural logarithm; logarithms to another base carry a subscript.

For $S \subseteq \mathbb{N}_{\geq 1}$, write

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

The set S has natural density one when $B_S(X) = o(X)$. Earlier natural-density results gave fixed-power descent. Allouche [2] proved $T_{\min}(n) \leq n^{\theta}$ for every fixed $\theta > 3/2 - (\log 2)/(\log 3)$;

Korec [5] improved the range to $\theta > \log_4 3$; and Inselmann [4] obtained every fixed exponent $\theta > 0$. At a different density level, Tao [7] proved descent below every function tending to infinity in logarithmic density. Concurrent preprints of Allikvere [1] and Mazur [6] obtain natural-density conclusions with arbitrary diverging thresholds through bridges built on Tao's framework. Their terminal-threshold class is more general than the theorem below. The threshold proved here is nevertheless eventually smaller than n^θ for every fixed $\theta > 0$.

Set

$$a_0 = \frac{\log_2 3}{2}, \qquad \delta_0 = \frac{\log(1/a_0)}{\log(2/a_0)} = 0.251245530155874\ldots. \tag{1.1}$$

The endpoint δ_0 is a supremum of the strict parameter ranges used in the proof; equality is not asserted.

Theorem 1.1

Fix $0 < \delta < \delta_0$.

1. The set of n satisfying

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

has natural density one. More quantitatively, for every

$$0 < \sigma < 1 - \frac{\delta}{\delta_0},$$

there are constants $c_{\delta,\sigma} > 0$ and $X_{\delta,\sigma}$ such that, for every $X \ge X_{\delta,\sigma}$,

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

2. There is a set A_δ of natural density one such that every sufficiently large $n \in A_\delta$ has an integer $0 \le k < 6.953 \log n$ for which

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

The constants in (1.3) are effective in the sense that they are obtained from the fixed parameters and finite maxima in the proof. They are not optimized, and no practical numerical value of $X_{\delta,\sigma}$ is claimed.

The restriction on σ comes from the shell exponent $1-\gamma$, where $\gamma = \omega \log(1/\chi)$ in (6.17); the corresponding parameter choice is made in (7.11)–(7.12).

Theorem 1.1 is an almost-all theorem. It neither proves descent for every initial value nor excludes exceptional cycles or divergent trajectories. Section 8.1 records the precise scope of the conclusion.

1.1. Why logarithmic blocks appear

Suppose a block contains M half-Collatz steps and s odd steps. If the additive terms are temporarily ignored, its multiplicative factor is

$$\frac{3^s}{2^M}.$$

For a typical parity word $s \approx M/2$, so this factor is approximately

$$\left(\frac{\sqrt{3}}{2}\right)^M.$$

Taking $M \approx \log_2 n$ therefore sends the scale n to approximately n^{a_0}, where $a_0 = (\log_2 3)/2 < 1$. One logarithmic block gives a fixed-power contraction. Repeating such blocks is what converts that contraction into the stretched-logarithmic threshold in (1.2).

The choice of exactly M steps also matches the arithmetic of the shell: its 2^M starting values form a complete residue system modulo 2^M, so Proposition 2.3 identifies them with all length-M parity words exactly once. Shorter blocks retain exact uniformity but give less contraction, while longer blocks no longer sample every parity word from one shell.

The difficulty is not the heuristic contraction. It is preserving a quantitative natural-density estimate when the endpoint of one block becomes the starting value of the next. The proof isolates this as a pullback problem. For a set S, define

$$\mathrm{Pull}(S) = \{n \geq 1 : T^{\lfloor \log_2 n \rfloor}(n) \in S\}. \tag{1.5}$$

The central technical result shows that a power-saving bound for the complement of S survives this operation with an asymptotically linear loss in its exponent.

1.2. Proof outline

The proof has six steps.

1. **One block in parity coordinates.** Parity words of length M are in bijection with the integers in a dyadic shell. The exact affine iterate expresses the block endpoint in terms of the parity word.
2. **Condition on the odd count.** After fixing the number s of odd steps, a parity word becomes a positive composition of a fixed total. Equality of block endpoints is reduced to a congruence modulo 3^s.
3. **Control endpoint multiplicities.** Words sharing one block endpoint form an endpoint fiber. Integer-lift injectivity and a quotient-order estimate convert congestion in its congruence class into a random variable measuring how much of the composition total is spent in its final parts. A Rényi bound at each fixed total controls its $(1+\theta)$-moment when the number of parts is near its mean, for every fixed $1/2 < \theta < 1$.
4. **Pull back a dense target.** Split endpoint fibers according to whether their normalized multiplicity exceeds a chosen exponential threshold; call those above it heavy. The target-density bound controls the ordinary part, and the Rényi estimate controls the heavy part. Balancing the two errors gives the small-D transport slope

$$\Lambda(\theta) = a_0 \frac{\theta}{1+\theta}. \tag{1.6}$$

Concretely, the ordinary and heavy contributions decay at rates $\psi(D) - v$ and θv; equating them gives the factor $\theta/(1+\theta)$, and the remaining factor a_0 is the small-D slope of $\psi(D)/\log 2$.
5. **Iterate endpoints only.** A terminal binomial-tail estimate supplies a fixed dense set on which one logarithmic block contracts n to at most $3n^r$, for any fixed $r > a_0$. The proof recursively pulls back this fixed contraction set; it does not track intermediate iterates inside a block. Requiring a shrinking envelope at every intermediate time would impose a stronger pathwise event not supplied by the endpoint pullback in Step 4. For an orbit-minimum theorem, endpoint iteration is the minimal recursion matching the conclusion.
6. **Choose the number of blocks.** Using $R \asymp \log\log n$ blocks gives endpoint size n^{r^R}. The density recurrence remains summable exactly when the block-growth parameter lies between two explicit scalar bounds. Letting $r \downarrow a_0$ and $\theta \uparrow 1$ gives (1.1).

This organization separates the two genuinely different tasks: proving a one-block transport theorem and iterating its endpoint map. Every estimate is stated in the form required by the next step.

1.3. Relation to other density results

Relative to the results above, the proof is independent of Tao's first-passage framework and uses no Diophantine approximation input. It trades the arbitrary diverging thresholds available through the concurrent bridge approaches for a prescribed stretched-logarithmic threshold, the counting-cutoff estimate (1.3), and the logarithmic-time witness (1.4). Its inputs are finite parity coding, compositions conditioned on their total, congruence-fiber injectivity, one Rényi estimate near the mean composition length, and a binomial tail bound.

The new quantitative mechanism is Theorem 5.3. Once its moment parameters are fixed, the pullback constants remain fixed while the recursively generated density exponent tends to zero.

1.4. Organization

Section 2 gives the density calculus and exact one-block coordinates. Sections 3 and 4 prove the Rényi estimate for compositions conditioned on their total and with part count near its mean. Section 5 turns that estimate into a quantitative pullback theorem. Section 6 proves the one-block contraction and endpoint bootstrap. Section 7 optimizes the parameters and proves Theorem 1.1. Section 8 records scope, formal verification supplementary to the manuscript proof, and declarations.

1.5. Terminology and recurring notation

We use the following terms throughout. A dyadic shell is an interval $[2^M, 2^{M+1})$; a logarithmic block starting there consists of M iterations, and its endpoint is the terminal value $T^M(n)$. A cohort is a set of parity words with a specified odd count and, when stated, a specified first odd position. *Fixed-total* means that the composition total N is held fixed, while *central* means that its number of parts lies within εN of the mean $(N+1)/2$, for the fixed ε in (4.6). An endpoint fiber consists of starts sharing one endpoint; in Section 5 a fiber is *heavy* precisely when it satisfies the threshold inequality (5.13), and is ordinary otherwise. *Transport* refers to passing a quantitative density estimate from a target set to its pullback under the block endpoint map.

For reference, the recurring notation is grouped by its role in the proof. The last column points to the defining display; symbols omitted from the table are local to a single argument.

After δ, and also σ for the quantitative assertion, are fixed, the auxiliary parameters $\theta, \varepsilon, r, \chi, \zeta, \omega$ are chosen once and remain fixed across shells. Only the recursively propagated density parameters C_j, D_j and the shell-dependent block count R_M vary during the iteration.

Stage	Recurring notation and role	Definition
Density and headline constants	$T_{\min}(n)$, least orbit value; $B_S(X)$, exceptional count; a_0, one-block exponent; δ_0, theorem endpoint	Abstract; Introduction; (1.1)
Fixed-total fibers	$\mathscr{A}_{\mathbf{k}}$, composition correction; $L_{s,N}$, composition count; $\nu_{s,N}$, residue law; $\mathcal{A}_{s,N}(\theta)$, normalized Rényi functional	(3.4), (3.6)–(3.7)
Central composition law	J_N, random number of parts; $p_{N,s}$, its mass function; $\varepsilon_0(\theta)$, admissible half-width; $\mathcal{I}_{N,\varepsilon}$, central index set	(4.1), (4.5)–(4.6)

Stage	Recurring notation and role	Definition
Shell transport	G_M^{cen}, central shell moment; $\psi(D)$, ordinary-fiber exponent; $r_\theta(D)$, balanced transport rate; $\Lambda(\theta)$, limiting transport slope	(5.3), (5.8), (5.12)
Endpoint iteration	F, logarithmic-block endpoint; η_r, odd-count margin; $\mathcal{E}_r$, contraction set; K_{blk}, block constant; $\mathcal{S}_j$, recursive good sets	(6.1)–(6.2), (6.7)

2. One logarithmic block

The proof begins on one dyadic shell, where a logarithmic block has a fixed length. Two translations are needed before its endpoints can be counted. First, shell estimates must be converted into estimates for all initial intervals. Second, a starting integer must be replaced by its finite parity word and the endpoint written exactly in those coordinates. Lemmas 2.1–2.2 handle the density and concentration estimates; Propositions 2.3–2.4 and Corollary 2.5 provide the endpoint coordinates used in Section 3.

2.1. From dyadic shells to initial intervals

For $C > 0$ and $0 < D \le 1$, call a set $S \subseteq \mathbb{N}_{\ge 1}$ (C, D)**-dense** if

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

Every such set has natural density one.

We need two elementary ways to pass from dyadic-shell estimates to (2.1).

Lemma 2.1 (dyadic summation)

1. Suppose $K > 0$, $0 < \gamma < \log 2$, and

$$\#\big(S^c \cap [2^M, 2^{M+1})\big) \le Ke^{-\gamma M}2^M \qquad (M \ge 0). \tag{2.2}$$

Then S is

$$\left(\frac{2K}{2e^{-\gamma}-1}, \frac{\gamma}{\log 2}\right)\text{-dense.} \tag{2.3}$$

2. Suppose $0 < \sigma \le 1$, $A, \gamma > 0$, and, for all sufficiently large M,

$$\#\big(E \cap [2^M, 2^{M+1})\big) \le Ae^{-\gamma(M+4)^\sigma}2^M. \tag{2.4}$$

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

$$\#(E \cap [1, X]) \le A'X \exp\left(-\gamma'(\log X)^\sigma\right) \qquad (X \ge X_0). \tag{2.5}$$

One may take

$$A' = 2A + 1, \qquad \gamma' = \min\left\{\frac{\gamma}{(4\log 2)^\sigma}, \frac{1}{2}\right\}. \tag{2.6}$$

Proof.

For the first assertion, let $2^J \le X < 2^{J+1}$. Summing (2.2) over $0 \le M \le J$ gives

$$\begin{aligned} B_S(X) &\le K \sum_{M=0}^{J} (2e^{-\gamma})^M \\ &\le \frac{K(2e^{-\gamma})^{J+1}}{2e^{-\gamma}-1} \le \frac{2K}{2e^{-\gamma}-1} X^{1-\gamma/\log 2}, \end{aligned}$$

which is (2.3).

For the second assertion, choose M_* so that (2.4) holds for $M \ge M_*$, and take

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

When $0 < \sigma < 1$, enlarge X_0 so that

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

Again write $2^J \le X < 2^{J+1}$. The shells with $M < \lfloor J/2 \rfloor$ contain fewer than

$$\sum_{M < \lfloor J/2 \rfloor} 2^M < 2^{\lfloor J/2 \rfloor} \le X^{1/2} \tag{2.7}$$

integers. On each remaining shell,

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

so their total contribution is at most

$$2AX \exp\left(-\gamma'(\log X)^\sigma\right). \tag{2.8}$$

The definition of X_0 gives $\gamma'(\log X)^\sigma \le \frac{1}{2}\log X$; for $\sigma = 1$, this follows from $\gamma' \le 1/2$. Thus the right side of (2.7) is at most

$$X \exp\left(-\gamma'(\log X)^\sigma\right).$$

Adding (2.7) and (2.8) proves (2.5). □

We shall repeatedly use three immediate consequences of (2.1):

- the intersection of (C_i, D_i)-dense sets, $i = 1, 2$, is $(C_1 + C_2, \min(D_1, D_2))$-dense;
- adding or deleting finitely many integers changes only the constant C;
- if $0 < D' \le D$, then every (C, D)-dense set is also (C, D')-dense.

The central-window argument also uses the following symmetric-binomial estimate.

Lemma 2.2 (symmetric-binomial tail)

If $X \sim \mathrm{Bin}(n, 1/2)$, $n \ge 1$, and $y \ge 0$, then

$$\Pr\left(\left|X - \frac{n}{2}\right| \ge y\right) \le 2\exp\left(-\frac{2y^2}{n}\right). \tag{2.9}$$

Proof.

For $\lambda > 0$,

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

The last inequality follows from $\cosh u \le e^{u^2/2}$: on $[0,\infty)$, the derivative of $u^2/2 - \log\cosh u$ is $u - \tanh u$, whose derivative is $\tanh^2 u \ge 0$. Markov's inequality with $\lambda = 4y/n$ gives the upper tail. Applying the same argument to $-X$ gives the lower tail. □

2.2. Parity words and exact endpoints

The finite parity-word coordinates are classical; see Terras [8] and Everett [3]. We include the short proofs needed here.

For $n \geq 1$, define the parity bits and their prefix sums by

$$p_i(n) = \mathbf{1}_{\{T^i(n)\ \text{odd}\}}, \qquad s_k(n) = \sum_{i=0}^{k-1} p_i(n). \tag{2.10}$$

Binomial concentration is useful only if parity words occur with exact, not heuristic, uniformity among shell starts. The next proposition supplies that identification.

Proposition 2.3 (parity-word bijection)

For every $M \geq 0$, the map

$$n \bmod 2^M \longmapsto \big(p_0(n), \ldots, p_{M-1}(n)\big) \tag{2.11}$$

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

$$\#\{n \in [2^M, 2^{M+1}) : s_M(n) = s\} = \binom{M}{s}. \tag{2.12}$$

Proof.

The first M parity bits depend only on $n \bmod 2^M$, because

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

Conversely, the parity bit at time zero and the residue of $T(n)$ determine n one power of two more precisely. On the even branch,

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

while on the odd branch,

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

where 3 is a unit modulo 2^{k+1}. Induction recovers $n \bmod 2^M$ from the parity word. The two finite sets have the same cardinality, so the map is bijective. Each residue modulo 2^M occurs once in $[2^M, 2^{M+1})$, and (2.12) follows. □

Define an integer correction by

$$c_0(n) = 0, \qquad c_{k+1}(n) = \begin{cases} c_k(n), & p_k(n) = 0, \\ 3c_k(n) + 2^k, & p_k(n) = 1. \end{cases} \tag{2.13}$$

The parity word specifies the multiplicative factor $3^{s_k}/2^k$, but the additive $+1$ terms still have to be retained. The next identity keeps that correction exact rather than replacing the Collatz block by its multiplicative heuristic.

Proposition 2.4 (affine iterate)

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

$$2^k T^k(n) = 3^{s_k(n)} n + c_k(n), \tag{2.14}$$

and

$$0 \le c_k(n) < 3^{s_k(n)} 2^k. \tag{2.15}$$

Proof.

Both statements follow by induction on k. An even step leaves the correction unchanged. At an odd step,

$$c_{k+1} = 3c_k + 2^k, \qquad s_{k+1} = s_k + 1.$$

If (2.15) holds at time k, then

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

because $1 < 3^{s_k+1}$. This proves (2.14)–(2.15) simultaneously. □

For the fiber argument we shall not need a sharp bound on every correction. The affine identity already gives the following uniform endpoint range on a fixed shell and odd-count cohort.

Corollary 2.5 (endpoint range)

If $n \in [2^M, 2^{M+1})$ and $s_M(n) = s$, then

$$0 < T^M(n) < 3^{s+1}. \tag{2.16}$$

Proof.

Equations (2.14)–(2.15) 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 fibers

Section 2 replaces every shell start by a parity word and expresses its block endpoint through an integer correction. Pulling back a target set requires a bound on how many words can share one endpoint. We condition on the number s of odd letters, separate the first odd position, and encode the remaining gaps by a positive composition of a fixed total. This turns the endpoint-collision problem into a congruence-fiber problem modulo 3^s.

Fix a parity word $w \in \{0,1\}^M$ with $s \ge 1$ odd letters. The first lemma gives the exact dictionary between the word and its composition.

Lemma 3.1 (parity words and compositions)

Let

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

be the odd positions of w. Put

$$u = i_1, \qquad k_j = i_{j+1} - i_j \quad (1 \le j < s), \qquad k_s = M - i_s. \tag{3.1}$$

Then $\mathbf{k} = (k_1, \ldots, k_s)$ is a positive composition of $N = M - u$. Conversely, $u \ge 0$ and a positive composition of $M - u$ into s parts recover a unique parity word through

$$i_m = u + k_1 + \cdots + k_{m-1}. \tag{3.2}$$

If c_w is the correction in (2.13), then

$$c_w = 2^u \sum_{m=1}^{s} 3^{s-m} 2^{k_1+\cdots+k_{m-1}}. \tag{3.3}$$

Proof.

The gaps in (3.1) are positive and sum to $M-u$. Formula (3.2) proves the inverse correspondence. Each odd step at position i_m contributes 2^{i_m} to the correction and is multiplied by 3 at each of the $s-m$ later odd steps. Hence

$$c_w = \sum_{m=1}^{s} 3^{s-m} 2^{i_m},$$

which becomes (3.3) after substituting (3.2). □

For a positive composition

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

put $K_0 = 0$, $K_j = k_1 + \cdots + k_j$, and define

$$\mathscr{A}_{\mathbf{k}} = \sum_{j=0}^{s-1} 3^{s-1-j} 2^{K_j}. \tag{3.4}$$

Thus (3.3) says $c_w = 2^u \mathscr{A}_{\mathbf{k}}$. Since $M = u+N$, with 2^{-N} in the following congruence denoting the inverse of 2^N in $\mathbb{Z}/3^s\mathbb{Z}$,

$$2^{-M} c_w = 2^{-N} \mathscr{A}_{\mathbf{k}} \pmod{3^s}. \tag{3.5}$$

There are

$$L_{s,N} = \binom{N-1}{s-1} \tag{3.6}$$

positive compositions of N into s parts. Let $\nu_{s,N}$ be the probability law of

$$2^{-N} \mathscr{A}_{\mathbf{k}} \pmod{3^s}$$

when $\mathbf{k}$ is chosen uniformly from these compositions. For $\theta > 0$, define

$$\mathcal{A}_{s,N}(\theta) = 3^{s\theta} \sum_{a \bmod 3^s} \nu_{s,N}(a)^{1+\theta}. \tag{3.7}$$

Because 2^{-N} is a unit modulo 3^s, multiplication by it only permutes residue classes. The same functional is therefore obtained from the fibers of $\mathscr{A}_{\mathbf{k}} \bmod 3^s$.

The Rényi functional in (3.7) measures congestion in these residue fibers. To control it by ordering integer representatives, we first need to know that two different compositions never give the same integer correction.

Lemma 3.2 (integer-lift injectivity)

For fixed s, N, the map

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

is injective.

Proof.

Dividing (3.4) by 2^N gives

$$2^{-N}\mathscr{A}_{\mathbf{k}} = 3^{s-1}2^{-N} + 3^{s-2}2^{-(N-k_1)} + \cdots + 2^{-k_s}. \tag{3.8}$$

Taking valuations in $\mathbb{Q}_2$, the summands have valuations

$$K_0 - N < K_1 - N < \cdots < K_{s-1} - N.$$

Suppose two compositions give the same value in (3.8). Their first summands are identical, so cancel them. The 2-adic valuation of the remaining sum is $K_1 - N$: its first coefficient is odd, and every later summand has strictly larger valuation. Equality of the two remaining sums therefore forces the two values of K_1, and hence of k_1, to agree. Cancel the now identical next summands and repeat. This recovers $k_2, \ldots, k_{s-1}$; the common total N then recovers k_s. □

Injectivity allows the values in each residue class to be ordered by their quotients after division by 3^s. Distinct quotients must spread out, which converts a large fiber into a large average correction.

Lemma 3.3 (quotient-order estimate)

For every $\theta > 0$,

$$\mathcal{A}_{s,N}(\theta) \le (1+\theta)\,\mathbb{E}_{s,N}\left[\left(\frac{3^s + \mathscr{A}_{\mathbf{k}}}{L_{s,N}}\right)^{\theta}\right], \tag{3.9}$$

where $\mathbb{E}_{s,N}$ denotes expectation under the uniform law on positive compositions of N into s parts.

Proof.

For $a \in \{0, \ldots, 3^s - 1\}$, let f_a be the number of compositions with $\mathscr{A}_{\mathbf{k}} \equiv a \pmod{3^s}$. By Lemma 3.2, the integer values 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 $q_i \ge i - 1$, so

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

The last inequality uses $3^{-s}(3^s + \mathscr{A}_i) = 1 + q_i + a/3^s \ge 1 + q_i$. Multiply by $3^{s\theta}$, sum over a, and divide by $L_{s,N}^{1+\theta}$. This is (3.9). □

Thus the arithmetic collision moment has been reduced to the expectation on the right of (3.9). Section 4 estimates that expectation when the composition length lies in the central range relevant to typical parity words.

4. The central fixed-total Rényi estimate

The quotient-order estimate leaves one probabilistic task: average the size of the correction over compositions whose number of parts is close to half the total. Independent separators provide the natural model for this average. They also make terminal composition blocks explicit, which is what permits the estimate for terminal composition parts below. The resulting moment is subcritical for $\theta < 1$ in a fixed central window.

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

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

For $0 < \theta < 1$, put

$$\alpha_\theta = \frac{2^\theta - 1}{3^\theta - 1}. \tag{4.2}$$

Strict concavity of x^θ at the midpoint $2 = (1+3)/2$ gives

$$\frac{1}{2} < \alpha_\theta. \tag{4.3}$$

We shall also need

$$\alpha_\theta < \frac{\log 2}{\log 3}. \tag{4.4}$$

Indeed, the function

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

is strictly increasing on $(1, \infty)$. After writing $z = \theta \log x$, its derivative has the sign of $e^z(z-1)+1$, whose derivative is $ze^z > 0$ and whose value at zero is zero. Comparing $x = 2$ and $x = 3$ proves (4.4).

For $1/2 < \theta < 1$, define the available central half-width

$$\varepsilon_0(\theta) = \alpha_\theta - \frac{1}{2} = \frac{2^\theta - 1}{3^\theta - 1} - \frac{1}{2} > 0 \tag{4.5}$$

and, for $0 < \varepsilon < \varepsilon_0(\theta)$, set

$$\mathcal{I}_{N,\varepsilon} = \left\{ s : \left| s - \frac{N+1}{2} \right| \le \varepsilon N \right\}. \tag{4.6}$$

The central half-width in (4.5) is chosen precisely so that the geometric series over terminal composition parts used in the proof has ratio below one. The next theorem is the only higher-moment input needed later.

Theorem 4.1 (central fixed-total estimate)

Fix $1/2 < \theta < 1$ and $0 < \varepsilon < \varepsilon_0(\theta)$. There is a constant $C_{\theta,\varepsilon} < \infty$ such that, for every $N \ge 1$,

$$\sum_{s \in \mathcal{I}_{N,\varepsilon}} p_{N,s} \mathcal{A}_{s,N}(\theta) \le C_{\theta,\varepsilon} N^{1+\theta}. \tag{4.7}$$

Proof.

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

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

We first assume $1 < s < N$. The extreme cases are deterministic and will be absorbed into the finite set of small pairs at the end of the proof. Under the uniform law on positive compositions of N into s parts, prescribe final parts $d_1, \dots, d_r$, with

$$D = d_1 + \cdots + d_r < N,$$

and require another part before them. The exact probability is

$$\frac{\binom{N-1-D}{s-1-r}}{\binom{N-1}{s-1}}. \tag{4.8}$$

Bound it by viewing the separators as independent Bernoulli(π) bits conditioned on having exactly $s-1$ successes. The prescribed terminal block requires r successes and $D-r$ failures, so its unconditioned probability is $\pi^r(1-\pi)^{D-r}$.

The conditioning event has probability at least $1/N$. To see this, let $X \sim \mathrm{Bin}(N-1, \pi)$. Its mean $s-1$ is an integer, and

$$\frac{\Pr(X = j+1)}{\Pr(X = j)} = \frac{N-1-j}{j+1}\frac{\pi}{1-\pi}.$$

This ratio is greater than one at $j = s-2$, less than one at $j = s-1$, and monotone in j. Thus $s-1$ is a mode. Since X has N possible values, its largest mass is at least $1/N$. If A is the event prescribing the terminal block, then $\Pr(A \mid X = s-1) \le \Pr(A)/\Pr(X = s-1) \le N \Pr(A)$. Consequently,

$$\Pr_{s,N}(d_1, \dots, d_r \text{ terminal}) \le N\pi^r(1-\pi)^{D-r}. \tag{4.9}$$

For a fixed composition, reverse its parts and put

$$U_r = k_{s-r+1} + \cdots + k_s, \qquad Z_{s,N} = \frac{3^s + \mathscr{A}_{\mathbf{k}}}{2^N}. \tag{4.10}$$

Reindexing (3.4) gives the exact identity

$$Z_{s,N} = \sum_{r=1}^{s-1} 3^{r-1}2^{-U_r} + 4 \cdot 3^{s-1}2^{-N}. \tag{4.11}$$

Choose a number π_* satisfying

$$\frac{1}{2} + \varepsilon < \pi_* < \alpha_\theta. \tag{4.12}$$

For all sufficiently large N, every $s \in \mathcal{I}_{N,\varepsilon}$ satisfies $1 < s < N$ and $\pi \le \pi_*$. Define

$$R_* = \frac{3^\theta \pi_*}{2^\theta - 1 + \pi_*}. \tag{4.13}$$

The inequality $\pi_* < \alpha_\theta$ is equivalent to $R_* < 1$.

Because $0 < \theta < 1$, subadditivity of x^θ, followed by (4.9), gives

$$\begin{aligned} \mathbb{E}_{s,N} \sum_{r=1}^{s-1} (3^{r-1}2^{-U_r})^\theta &\le N3^{-\theta} \sum_{r \ge 1} \left[3^\theta \pi \sum_{d \ge 1} 2^{-\theta d}(1-\pi)^{d-1} \right]^r \\ &\le N3^{-\theta} \frac{R_*}{1-R_*}. \end{aligned} \tag{4.14}$$

Indeed, the expression in square brackets is

$$\frac{3^\theta \pi}{2^\theta - 1 + \pi},$$

which is increasing in π and is therefore at most R_*. For the actual term indexed by r, the terminal total satisfies $U_r < N$; extending the sum in (4.14) to all positive $d_1, \ldots, d_r$ only enlarges its right-hand side.

By (4.4) and (4.12), there is $N_0 = N_0(\theta, \varepsilon)$ such that

$$\frac{s}{N} \le \frac{\log 2}{\log 3}$$

whenever $N \ge N_0$ and $s \in \mathcal{I}_{N,\varepsilon}$. The final term in (4.11) is then at most 4/3. Equations (4.11) and (4.14) consequently give

$$\mathbb{E}_{s,N} Z_{s,N}^\theta \le K_{\theta,\varepsilon} N \qquad (N \ge N_0,\ s \in \mathcal{I}_{N,\varepsilon}) \tag{4.15}$$

for a fixed constant $K_{\theta,\varepsilon}$. Enlarge this constant by the maximum of

$$N^{-1} \mathbb{E}_{s,N} Z_{s,N}^\theta$$

over the finite set

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

If this set is empty, take the maximum to be zero. Then (4.15) holds for every $N \ge 1$ and every central s.

Since $L_{s,N} = 2^{N-1} p_{N,s}$, one has

$$\left(\frac{3^s + \mathscr{A}_{\mathbf{k}}}{L_{s,N}} \right)^\theta = 2^\theta p_{N,s}^{-\theta} Z_{s,N}^\theta. \tag{4.16}$$

Lemma 3.3 and (4.15) imply

$$\mathcal{A}_{s,N}(\theta) \le (1+\theta) 2^\theta K_{\theta,\varepsilon} N p_{N,s}^{-\theta}. \tag{4.17}$$

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

$$\begin{aligned} \sum_{s \in \mathcal{I}_{N,\varepsilon}} p_{N,s} \mathcal{A}_{s,N}(\theta) &\le (1+\theta) 2^\theta K_{\theta,\varepsilon} N \sum_{s=1}^N p_{N,s}^{1-\theta} \\ &\le (1+\theta) 2^\theta K_{\theta,\varepsilon} N^{1+\theta}. \end{aligned} \tag{4.18}$$

For the last step,

$$\frac{1}{N} \sum_{s=1}^N p_{N,s}^{1-\theta} \le \left(\frac{1}{N} \sum_{s=1}^N p_{N,s} \right)^{1-\theta} = N^{\theta - 1}.$$

Thus (4.7) holds with

$$C_{\theta,\varepsilon} = (1+\theta) 2^\theta K_{\theta,\varepsilon}.$$

□

Theorem 4.1 completes the fixed-total part of the argument. Its bound concerns one composition total N; the next section must transfer it back to all Collatz endpoint fibers in a shell and then use it against an arbitrary dense target.

5. From endpoint moments to density transport

Theorem 4.1 controls a fixed total and a central range of composition lengths. A Collatz shell mixes the possible first odd positions and also contains noncentral cohorts. Section 5.1 recombines the central cohorts and proves that the discarded source mass is exponentially small. Section 5.2 then splits the resulting endpoint fibers into ordinary and heavy parts. The output is a one-block theorem describing exactly how a quantitative density exponent changes under pullback.

The central half-width ε is fixed before the target-density exponent D. This order matters because the pullback will be iterated at exponents tending to zero. If the window changed with D, both the discarded-cohort rate and the moment constant would vary from stage to stage; the theorem below would no longer provide one uniform recurrence. With ε fixed, those constants remain fixed and the small-D transport rate can be absorbed uniformly.

5.1. The central endpoint fibers

Fix M and $1 \le s \le M$. For a first-odd position u, put $N = M - u$, and let $f_{M,s,u}(y)$ be the number of starts in $[2^M, 2^{M+1})$ whose parity word has weight s, first odd position u, and endpoint $T^M(n) = y$.

If two starts in this cohort have the same endpoint, (2.14) implies that their corrections are congruent modulo 3^s. By (3.5), every endpoint fiber therefore lies inside one residue fiber counted by $\nu_{s,N}$. If $c_{s,N}(a) = L_{s,N}\nu_{s,N}(a)$, then, for every $q \ge 1$,

$$\sum_{y} f_{M,s,u}(y)^q \le \sum_{a \bmod 3^s} c_{s,N}(a)^q. \tag{5.1}$$

Indeed, inside one correction class the sum of the endpoint multiplicities is at most $c_{s,N}(a)$, and the q-th power of that sum dominates the sum of the q-th powers.

Fix $1/2 < \theta < 1$ and $0 < \varepsilon < \varepsilon_0(\theta)$. Retain a cohort u only when

$$s \in \mathcal{I}_{M-u,\varepsilon}, \tag{5.2}$$

and let $f^{\mathrm{cen}}_{M,s}(y)$ be the sum of $f_{M,s,u}(y)$ over the retained cohorts. Set $f^{\mathrm{cen}}_{M,0}(y) = 0$, and define

$$G^{\mathrm{cen}}_M(\theta,\varepsilon) = 2^{-M} \sum_{s,y} f^{\mathrm{cen}}_{M,s}(y) \left(\frac{3^{s+1} f^{\mathrm{cen}}_{M,s}(y)}{\binom{M}{s}} \right)^{\theta}. \tag{5.3}$$

The first transfer step is to show that the fixed-total moment remains polynomial after the retained first-odd-position cohorts are recombined.

Lemma 5.1 (central endpoint moment)

There is $C'_{\theta,\varepsilon} < \infty$ such that

$$G^{\mathrm{cen}}_M(\theta,\varepsilon) \le C'_{\theta,\varepsilon}(M+1)^{1+\theta} \qquad (M \ge 1). \tag{5.4}$$

Proof.

For fixed M, s, put

$$\lambda_u = \frac{L_{s,M-u}}{\binom{M}{s}}.$$

The first-odd-position cohorts partition the weight-s words, so the λ_u sum to one over all possible u, and to at most one over the retained u. Adjoin a zero cohort carrying the missing weight. The weights then sum to one, so weighted convexity and (5.1) give

$$\left(\sum_{u \text{ retained}} f_{M,s,u}(y) \right)^{1+\theta} \leq \sum_{u \text{ retained}} \lambda_u^{-\theta} f_{M,s,u}(y)^{1+\theta}.$$

After summing over y, (5.1) gives

$$\sum_y f_{M,s}^{\mathrm{cen}}(y)^{1+\theta} \leq \binom{M}{s}^{\theta} 3^{-s\theta} \sum_{\substack{u \text{ retained} \\ N=M-u}} L_{s,N} \mathcal{A}_{s,N}(\theta). \tag{5.5}$$

Substitute this into (5.3). Since

$$2^{-M} L_{s,M-u} = 2^{-u-1} p_{M-u,s},$$

Theorem 4.1 yields

$$\begin{aligned} G_M^{\mathrm{cen}}(\theta, \varepsilon) &\leq 3^{\theta} \sum_{u=0}^{M-1} 2^{-u-1} \sum_{s \in \mathcal{I}_{M-u,\varepsilon}} p_{M-u,s} \mathcal{A}_{s,M-u}(\theta) \\ &\leq 3^{\theta} C_{\theta,\varepsilon} \sum_{u=0}^{M-1} 2^{-u-1} (M-u)^{1+\theta} \\ &\leq 3^{\theta} C_{\theta,\varepsilon} (M+1)^{1+\theta}. \end{aligned}$$

This is (5.4). □

The moment estimate applies only to the central cohorts retained in (5.2). The next lemma shows that their complement is too small to affect the eventual pullback exponent.

Lemma 5.2 (discarded cohorts)

There are $C_\varepsilon, c_\varepsilon > 0$, depending only on ε, such that the normalized source mass outside the retained cohorts is at most

$$C_\varepsilon e^{-c_\varepsilon M}. \tag{5.6}$$

Proof.

At fixed u, with $N = M - u$, the normalized mass of all compositions is 2^{-u-1}, and their number of parts has law J_N from (4.1). For $N \geq 2$, Lemma 2.2 gives

$$\Pr\left(\left| J_N - \frac{N+1}{2} \right| > \varepsilon N \right) \leq 2e^{-2\varepsilon^2 N}. \tag{5.7}$$

The all-even word has mass 2^{-M}. Hence the discarded mass is at most

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

Choose

$$0 < c_\varepsilon < \min(2\varepsilon^2, \log 2).$$

Each summand is bounded by

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

whose sum over u is a fixed geometric series. Also $2^{-M} \leq e^{-c_\varepsilon M}$. This proves (5.6). □

5.2. Pulling back a dense target

Define

$$\psi(D) = -\log\left(\frac{1+3^{-D}}{2}\right), \qquad r_\theta(D) = \frac{\theta}{1+\theta}\psi(D). \tag{5.8}$$

Direct differentiation gives

$$\psi'(D) = \frac{\log 3}{3^D+1}, \qquad \psi''(D) = -\frac{(\log 3)^2\, 3^D}{(3^D+1)^2} < 0. \tag{5.9}$$

Thus ψ is increasing and concave on $[0,1]$, with

$$\psi(0) = 0, \qquad \psi(D) \geq D\log(3/2) \quad (0 \leq D \leq 1). \tag{5.10}$$

There are now three errors to balance. Target density controls ordinary endpoint fibers, Lemma 5.1 controls heavy fibers, and Lemma 5.2 controls the discarded cohorts. Choosing the heavy-fiber threshold so that the first two errors match produces the rate $r_\theta(D)$.

Theorem 5.3 (central endpoint pullback)

Fix $1/2 < \theta < 1$ and $0 < \varepsilon < \varepsilon_0(\theta)$. There is $D_{\theta,\varepsilon} > 0$ such that, for every $0 < \zeta < 1$, there is $K_{\theta,\varepsilon,\zeta} > 0$ with the following property: whenever $0 < D \leq D_{\theta,\varepsilon}$ and S is (C,D)-dense, $\mathrm{Pull}(S)$ is

$$\left(K_{\theta,\varepsilon,\zeta}(C+1)D^{-(1+\theta)}, \frac{\zeta r_\theta(D)}{\log 2}\right)\text{-dense}. \tag{5.11}$$

All constants in (5.11) are independent of C, D, S, the dyadic shell, and every later use of the theorem.

Consequently, if

$$0 < \chi < \Lambda(\theta) := a_0\frac{\theta}{1+\theta}, \tag{5.12}$$

then one may choose $0 < \zeta < 1$ and a further cutoff $D_{\mathrm{lin}} = D_{\mathrm{lin}}(\theta,\varepsilon,\zeta,\chi) > 0$ so that the density exponent in (5.11) is at least χD whenever $0 < D \leq D_{\mathrm{lin}}$.

Thus, after $\theta, \varepsilon, \chi$ and a compatible ζ are fixed, $K_{\theta,\varepsilon,\zeta}$ and D_{lin} do not change when later iterations send D to zero.

Proof.

Let S be (C,D)-dense, and fix a shell $[2^M, 2^{M+1})$. The shell $M = 0$ contains one integer and is absorbed by the final density constant, so assume $M \geq 1$. For $v > 0$, call a retained central endpoint fiber heavy when

$$\frac{3^{s+1} f^{\mathrm{cen}}_{M,s}(y)}{\binom{M}{s}} > e^{vM}. \tag{5.13}$$

Lemma 5.1 and Markov's inequality bound the normalized source mass in heavy fibers by

$$C'_{\theta,\varepsilon}(M+1)^{1+\theta}e^{-\theta v M}. \tag{5.14}$$

On every remaining central fiber,

$$f^{\mathrm{cen}}_{M,s}(y) \leq e^{vM}\binom{M}{s}3^{-(s+1)}. \tag{5.15}$$

Corollary 2.5 shows that every endpoint at parity level s lies below 3^{s+1}. Since S is (C,D)-dense, there are at most $C3^{(s+1)(1-D)}$ bad possible endpoints at this level. Therefore the ordinary-fiber contribution is at most

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

The sum is exact:

$$2^{-M}\sum_{s=0}^{M}\binom{M}{s}3^{-D(s+1)} = 3^{-D}\left(\frac{1+3^{-D}}{2}\right)^{M} = 3^{-D}e^{-\psi(D)M}. \tag{5.16}$$

Hence the ordinary contribution is

$$C3^{-D}e^{-(\psi(D)-v)M}. \tag{5.17}$$

Choose

$$v = \frac{\psi(D)}{1+\theta}. \tag{5.18}$$

Then

$$\theta v = \psi(D) - v = r_\theta(D).$$

Combining (5.6), (5.14), and (5.17), and using $3^{-D} \le 1$, gives

$$\frac{\#\{n \in [2^M, 2^{M+1}) : T^M(n) \notin S\}}{2^M} \le C_\varepsilon e^{-c_\varepsilon M} + K_{\theta,\varepsilon}(C+1)(M+1)^{1+\theta}e^{-r_\theta(D)M}. \tag{5.19}$$

Choose $D_{\theta,\varepsilon} \in (0,1]$ so that

$$r_\theta(D_{\theta,\varepsilon}) \le \min\{c_\varepsilon, \log(4/3)\}. \tag{5.20}$$

This choice is possible because r_θ is continuous, increasing, and vanishes at zero. It depends only on θ, ε. For $0 < D \le D_{\theta,\varepsilon}$, the first term in (5.19) is bounded by a constant multiple of $e^{-r_\theta(D)M}$. Since $C+1 \ge 1$ and $D^{-(1+\theta)} \ge 1$, this term can be absorbed into the same final constant as the central-fiber contribution.

We next absorb the polynomial factor without changing its dependence on D. For $B \ge 0$, $0 < u \le 1$, and $M \ge 1$,

$$(M+1)^B e^{-uM} \le K_{B,\zeta}u^{-B}e^{-\zeta uM}, \tag{5.21}$$

where $K_{B,\zeta}$ is independent of u, M. Indeed,

$$(M+1)^B u^B e^{-(1-\zeta)uM} \le (2uM)^B e^{-(1-\zeta)uM},$$

and the right side has a finite supremum as a function of $uM \ge 0$. Apply (5.21) with $B = 1+\theta$ and $u = r_\theta(D)$. By (5.10),

$$r_\theta(D) \ge \frac{\theta\log(3/2)}{1+\theta}D. \tag{5.22}$$

Thus (5.19) becomes

$$\frac{\#\{n \in [2^M, 2^{M+1}) : T^M(n) \notin S\}}{2^M} \le K_{\theta,\varepsilon,\zeta}(C+1)D^{-(1+\theta)}e^{-\zeta r_\theta(D)M}. \tag{5.23}$$

The fixed-rate part of Lemma 2.1 applies because, by (5.20), $\zeta r_\theta(D) < \log 2$. Its geometric-series factor is at most

$$\frac{2}{2e^{-\zeta r_\theta(D)} - 1} \le 4.$$

This proves (5.11), after enlarging $K_{\theta,\varepsilon,\zeta}$.

Finally fix $\chi < \Lambda(\theta)$ and choose $\zeta < 1$ so that $\chi < \zeta\Lambda(\theta)$. Since ψ' is decreasing and $\psi(0) = 0$, concavity gives $\psi(D) \ge D\psi'(D)$. Hence (5.9) gives

$$\frac{\zeta r_\theta(D)}{D\log 2} \ge \frac{\zeta\theta\log 3}{(1+\theta)\log 2\,(3^D+1)}. \tag{5.24}$$

The right side tends to $\zeta\Lambda(\theta) > \chi$ as $D \downarrow 0$. More explicitly, (5.24) is at least χ whenever

$$0 < D \le \frac{\log\left(\dfrac{\zeta\theta\log 3}{(1+\theta)\chi\log 2} - 1\right)}{\log 3}. \tag{5.25}$$

The quantity inside the logarithm is greater than one by the choice of ζ. Define D_{lin} to be the minimum of this bound, 1, and $D_{\theta,\varepsilon}$. This proves (5.12). □

This is the interface needed for iteration: for sufficiently small D, a pullback multiplies the density exponent by any fixed $\chi < \Lambda(\theta)$, while its density constant grows only polynomially in D^{-1}.

6. Endpoint iteration

The pullback theorem controls density from one block to the next, but it does not itself say that a block contracts. We therefore need a fixed dense seed set on which the endpoint map decreases scale. Intersecting that seed with successive pullbacks ensures both that the current block contracts and that its endpoint remains eligible for the remaining blocks. Lemma 6.1 constructs the seed set; Proposition 6.2 combines the resulting orbit recursion with the density recursion.

6.1. A dense one-block contraction

Define

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

Fix $a_0 < r < 1$, and put

$$\eta_r = \frac{r-a_0}{\log_2 3}, \qquad \mathcal{E}_r = \{n \ge 1 : F(n) \le 3n^r\}. \tag{6.2}$$

The heuristic $s \approx M/2$ is not enough for recursion: we need one fixed contraction set whose complement has a quantitative shell bound. Exact parity uniformity from Proposition 2.3 and the binomial tail in Lemma 2.2 provide it for every fixed $r > a_0$.

Lemma 6.1 (fixed contraction set)

There are constants $C_r, D_r > 0$ such that $\mathcal{E}_r$ is (C_r, D_r)-dense.

Proof.

Fix a shell $[2^M, 2^{M+1})$, and let s be the number of odd steps among the first M iterates. The case $M = 0$ is immediate. For $M \ge 1$, if

$$s \le \frac{M}{2} + \eta_r M, \tag{6.3}$$

then Corollary 2.5 gives

$$F(n) < 3^{s+1} \le 3 \cdot 2^{(a_0+\eta_r\log_2 3)M} = 3\cdot 2^{rM} \le 3n^r. \tag{6.4}$$

By Proposition 2.3, the parity words are uniform on $\{0,1\}^M$. Lemma 2.2 therefore gives

$$\frac{\#(\mathcal{E}_r^c \cap [2^M, 2^{M+1}))}{2^M} \le 2e^{-2\eta_r^2 M}. \tag{6.5}$$

Choose

$$0 < \gamma_r < \min\{2\eta_r^2, \log 2\}.$$

Then (6.5) implies (2.2) with $K = 2$ and $\gamma = \gamma_r$. Lemma 2.1 proves the claim, for example with

$$C_r = \frac{4}{2e^{-\gamma_r} - 1}, \qquad D_r = \frac{\gamma_r}{\log 2}. \tag{6.6}$$

□

Only the terminal odd count is used here. No control of intermediate iterates inside a block is needed for an orbit-minimum theorem.

6.2. The recursive endpoint sets

Set $K_{\mathrm{blk}} = 3$. With r fixed as above, define

$$\mathcal{S}_1 = \mathcal{E}_r, \qquad \mathcal{S}_{j+1} = \mathcal{E}_r \cap \mathrm{Pull}(\mathcal{S}_j). \tag{6.7}$$

If $n \in \mathcal{S}_R$, put

$$n_0 = n, \qquad n_{i+1} = F(n_i) \quad (0 \le i < R). \tag{6.8}$$

Induction in (6.7) gives $n_i \in \mathcal{E}_r$ for $0 \le i < R$. Consequently

$$n_i \le K_{\mathrm{blk}}^{1+r+\cdots+r^{i-1}} n^{r^i} \le K_{\mathrm{blk}}^{1/(1-r)} n^{r^i}. \tag{6.9}$$

Each n_i is an actual Collatz iterate of n. More precisely, if

$$\tau_R = \sum_{i=0}^{R-1} \lfloor \log_2 n_i \rfloor, \tag{6.10}$$

then $n_R = T^{\tau_R}(n)$, and (6.9) gives

$$\tau_R \le \frac{\log n}{(1-r)\log 2} + \frac{R \log K_{\mathrm{blk}}}{(1-r)\log 2}. \tag{6.11}$$

Equations (6.7)–(6.11) settle the deterministic part: membership in $\mathcal{S}_R$ gives an actual Collatz iterate with the required size and clock. It remains to show that R may grow like $\log\log n$ without the exceptional set becoming too large. The next proposition isolates that calculation from the specific pullback theorem.

Proposition 6.2 (endpoint bootstrap)

Fix $a_0 < r < 1$. Suppose there are constants

$$0 < \chi < a_0, \qquad B \ge 0, \qquad K_{\mathrm{L}}, D_c > 0 \tag{6.12}$$

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

$$(C, D) \longmapsto \big(K_{\mathrm{L}}(C+1)D^{-B}, \chi D\big). \tag{6.13}$$

For every $\omega > 0$ satisfying

$$\omega \log(1/\chi) < 1, \tag{6.14}$$

define, on the shell $[2^M, 2^{M+1})$,

$$R_M = \lceil \omega \log(M+4) \rceil. \tag{6.15}$$

Put

$$\mathcal{S}^{(\omega)} = \bigcup_{M \ge 0} \left(\mathcal{S}_{R_M} \cap [2^M, 2^{M+1})\right).$$

Then the shell exceptional proportion outside $\mathcal{S}_{R_M}$ is at most

$$2\exp\left(-c(M+4)^{1-\gamma}\right) \tag{6.16}$$

for all sufficiently large M, where $c > 0$ and

$$\gamma = \omega \log(1/\chi) < 1. \tag{6.17}$$

The set $\mathcal{S}^{(\omega)}$ has natural density one. On this set,

$$\log n_{R_M} \le \frac{\log K_{\mathrm{blk}}}{1-r} + (M+1)\log 2\,(M+4)^{-\omega\log(1/r)}, \tag{6.18}$$

and the witnessing time satisfies (6.11) with $R = R_M$.

In particular, every

$$0 < \delta < \frac{\log(1/r)}{\log(1/\chi)} \tag{6.19}$$

can be realized by a choice of ω for which

$$n_{R_M} \le \exp((\log n)^{1-\delta}) \tag{6.20}$$

on all sufficiently large retained shells.

Proof.

Choose

$$0 < D_1 \le \min\{D_r, D_c, 1\}.$$

After weakening the exponent in Lemma 6.1 to D_1, put $C_1 = C_r$. Then $\mathcal{E}_r$ is (C_1, D_1)-dense. Inductively, the sets in (6.7) are (C_j, D_j)-dense with

$$D_{j+1} = \chi D_j, \qquad C_{j+1} \le K_{\mathrm{L}}(C_j+1)D_j^{-B} + C_r. \tag{6.21}$$

Indeed, $D_j \le D_1 \le D_c$, so pullback is applicable, and $\chi D_j \le D_1 \le D_r$, so intersection with $\mathcal{E}_r$ preserves the exponent χD_j. Thus

$$D_j = D_1 \chi^{j-1}. \tag{6.22}$$

Let

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

Since $D_j \le 1$, (6.21) gives

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

Taking logarithms and using (6.22),

$$\begin{aligned}\log(C_R+2) \le{}& \log(C_1+2) + (R-1)\big(\log A + B\log D_1^{-1}\big)\\ &+ \frac{B}{2}(R-1)(R-2)\log(\chi^{-1}).\end{aligned} \tag{6.23}$$

In particular,

$$\log(C_R+2) = O(R^2). \tag{6.24}$$

From (6.15) and (6.22),

$$D_{R_M} \ge D_1(M+4)^{-\omega\log(1/\chi)}. \tag{6.25}$$

The (C_{R_M}, D_{R_M})-density bound implies that the exceptional proportion on shell M is at most

$$2C_{R_M}e^{-(\log 2)D_{R_M}M}. \tag{6.26}$$

Equations (6.14), (6.24), and (6.25) give

$$\log(2C_{R_M}) - (\log 2)D_{R_M}M \le O((\log M)^2) - c_1 M^{1-\gamma}. \tag{6.27}$$

After decreasing the positive constant and increasing the fixed starting shell, this proves (6.16). Lemma 2.1, part 2, then shows that the union of the retained shell sets has natural density one.

Equation (6.18) follows from (6.9), because

$$r^{R_M} \le (M+4)^{-\omega\log(1/r)} \quad \text{and} \quad \log n < (M+1)\log 2.$$

If

$$\delta < \omega\log(1/r), \tag{6.28}$$

then the right side of (6.18) is eventually at most $(\log n)^{1-\delta}$. Indeed, after division by that quantity, its second term is

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

and its fixed first term also tends to zero. Thus (6.20) holds. Conditions (6.14) and (6.28) admit a common ω exactly when (6.19) holds. □

Proposition 6.2 reduces the final theorem to a scalar compatibility condition: the endpoint must contract at rate r, while the density exponent survives each pullback at rate χ. Section 7 chooses these parameters using the slope supplied by Theorem 5.3.

7. Proof of Theorem 1.1

The arithmetic and iteration arguments are now complete. Theorem 5.3 permits every pullback slope below $\Lambda(\theta)$, while Proposition 6.2 converts such a slope and a contraction exponent r into a stretched-logarithmic descent exponent. What remains is to choose the parameters with strict margins, pass to the limiting endpoint, and retain the exceptional-count and time bounds.

For $1/2 < \theta < 1$, define

$$\Lambda(\theta) = a_0\frac{\theta}{1+\theta}, \qquad \delta_*(\theta) = \frac{\log(1/a_0)}{\log(1/\Lambda(\theta))}. \tag{7.1}$$

The first number is the limiting transport slope from Theorem 5.3; the second is the endpoint exponent supplied by Proposition 6.2 at that slope.

The limiting value of this scalar expression determines the range stated in Theorem 1.1.

Lemma 7.1 (endpoint limit)

The function δ_* is strictly increasing on $(1/2, 1)$, and

$$\lim_{\theta\uparrow 1}\delta_*(\theta) = \frac{\log(1/a_0)}{\log(2/a_0)} = \delta_0. \tag{7.2}$$

Proof.

The function $\theta \mapsto \theta/(1+\theta)$ is strictly increasing. Hence $\Lambda(\theta)$ is strictly increasing and lies in $(0,1)$; therefore $\log(1/\Lambda(\theta))$ is positive and strictly decreasing. This proves the monotonicity of δ_*. Since $\Lambda(\theta) \to a_0/2$ as $\theta \uparrow 1$, continuity gives (7.2). □

The strict endpoint reflects a boundary of the present moment argument. As $\theta \uparrow 1$, the admissible half-width $\varepsilon_0(\theta)$ in (4.5) tends to zero, and at $\theta = 1$ the geometric ratio in (4.14) reaches one at the central value $\pi = 1/2$. Thus Theorem 4.1 supplies neither a fixed central window nor a subcritical estimate at the endpoint. This does not rule out attaining or exceeding δ_0 by a different estimate.

We now choose the parameters in the order required by their dependencies: first the moment order and central window, then the contraction and pullback rates, and finally the number of blocks.

Proof of Theorem 1.1

Fix $0 < \delta < \delta_0$. By Lemma 7.1, choose

$$\frac{1}{2} < \theta < 1 \quad \text{such that} \quad \delta < \delta_*(\theta). \tag{7.3}$$

Equation (4.5) gives $\varepsilon_0(\theta) > 0$, so choose once and for all

$$0 < \varepsilon < \varepsilon_0(\theta). \tag{7.4}$$

By continuity, there are strict interior values

$$a_0 < r < 1, \qquad 0 < \chi < \Lambda(\theta) \tag{7.5}$$

such that

$$\delta < \frac{\log(1/r)}{\log(1/\chi)}. \tag{7.6}$$

Choose ζ and the cutoff D_{lin} from the consequence of Theorem 5.3. After weakening the resulting pullback exponent to χD, that theorem supplies (6.13) with

$$B = 1 + \theta, \qquad K_{\text{L}} = K_{\theta,\varepsilon,\zeta}, \qquad D_c = D_{\text{lin}}. \tag{7.7}$$

Choose

$$\frac{\delta}{\log(1/r)} < \omega < \frac{1}{\log(1/\chi)}. \tag{7.8}$$

Proposition 6.2 now gives a natural-density-one union of shell sets on which (6.20) holds. Since $n_{R_M} = T^{\tau_{R_M}}(n)$, this proves (1.2).

We next prove the quantitative assertion. Fix

$$0 < \sigma < 1 - \frac{\delta}{\delta_0}. \tag{7.9}$$

Equivalently, $\delta/(1-\sigma) < \delta_0$. Repeat the choices (7.3)–(7.7), now taking the strict interior values r, χ so that

$$\frac{\delta}{1-\sigma} < \frac{\log(1/r)}{\log(1/\chi)}. \tag{7.10}$$

Set

$$\omega = \frac{1-\sigma}{\log(1/\chi)}. \tag{7.11}$$

Then

$$\omega \log(1/\chi) = 1 - \sigma < 1, \qquad \delta < \omega \log(1/r). \tag{7.12}$$

Thus (6.16) bounds the shell exceptional proportion by

$$2\exp(-c_0(M+4)^{\sigma}) \tag{7.13}$$

for some $c_0 > 0$ and all sufficiently large M. Apply the second part of Lemma 2.1 with $A = 2$. Its explicit value $A' = 2A+1$ is 5, so, after increasing only $X_{\delta,\sigma}$ past the finitely many startup shells, there is $c_{\delta,\sigma} > 0$ for which

$$\#\left\{1 \le n \le X : T_{\min}(n) > \exp((\log n)^{1-\delta})\right\} \le 5X \exp\left(-c_{\delta,\sigma}(\log X)^{\sigma}\right) \tag{7.14}$$

for every $X \ge X_{\delta,\sigma}$. This is (1.3).

It remains to prove the advertised clock. First,

$$\frac{1}{(1-a_0)\log 2} = \frac{2}{\log(4/3)} < \frac{6953}{1000}. \tag{7.15}$$

For a rational certificate of the strict inequality, put $x = 1/7$ in

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

The first two positive terms give

$$\log \frac{4}{3} > \frac{296}{1029},$$

and

$$\frac{1029}{148} < \frac{6953}{1000}$$

because 1,029,000 < 1,029,044.

Choose θ, χ so that

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

with $1/2 < \theta < 1$ and $0 < \chi < \Lambda(\theta)$. Choose $\varepsilon, \zeta, D_{\mathrm{lin}}$ as in (7.4) and (7.7), so that the pullback hypothesis of Proposition 6.2 holds. Then take $r > a_0$ sufficiently close to a_0 that both (7.6) and

$$\frac{1}{(1-r)\log 2} < 6.953 \tag{7.16}$$

hold. With ω chosen as in (7.8), one has $R_M = O(\log M) = O(\log\log n)$. Equation (6.11) therefore gives

$$\tau_{R_M} \le \frac{\log n}{(1-r)\log 2} + O_r(\log\log n) < 6.953 \log n \tag{7.17}$$

for all sufficiently large retained n. Together with (6.20), this proves (1.4). The retained shell union has density one by Proposition 6.2, through Lemma 2.1, part 2, and is the required set A_δ. □

8. Scope and reproducibility

The proof is complete at this point. We finish by separating the theorem from stronger Collatz conclusions that it does not imply and by recording the formal artifact independently of the manuscript proof.

8.1. Scope and comparison

Theorem 1.1 is compatible with exceptional divergent trajectories and with hypothetical nontrivial cycles. If an orbit enters a fixed cycle with minimum m, then its orbit minimum is at most m, so all sufficiently large points in that basin already satisfy (1.2). The theorem therefore gives no cycle exclusion and no density estimate for such basins.

Tao [7] proves the stronger terminal-threshold statement with an arbitrary function tending to infinity in logarithmic density. The concurrent preprints [1, 6] obtain arbitrary-diverging-threshold statements in natural density through bridges built on Tao's framework. Theorem 1.1 has a less general terminal threshold; its separate features are the self-contained fixed-total endpoint argument, the prescribed stretched-logarithmic family, the counting-cutoff estimate (1.3), and the explicit half-map clock.

The fixed-power statements preceding this paper do not by themselves give the prescribed rate in (1.2). A density-one diagonalization over all fixed powers can produce some unspecified exponent $\varepsilon(n) \to 0$, but its decay depends on untracked onset thresholds. It does not imply $\varepsilon(n) = (\log n)^{-\delta}$, nor does it give (1.3).

8.2. Formal verification

The principal theorem chain has a supplementary Lean 4 formalization at

https://github.com/shaikidris/CET/tree/v2.0.1

with software DOI `10.5281/zenodo.21797535`. The frozen development uses Lean `v4.15.0` and Mathlib commit `9837ca9d65d9de6fad1ef4381750ca688774e608`. Its referee-facing Main Theorem, exceptional-count theorem, and timed theorem are respectively

```
collatz_central_renyi_endpoint_natural_density_descent
collatz_central_renyi_endpoint_exceptional_count_at_exponent
collatz_central_renyi_endpoint_natural_density_descent_timed
```

in namespace `CollatzEndpointTransport.QuantitativeCollatzMain`. The formal proof uses the stronger all-prefix initial-window estimate at its one-block seed; Lemma 6.1 uses only the terminal consequence needed by the endpoint theorem. The three declarations state exactly the conclusions (1.2), (1.3), and (1.4); only the intermediate route at the seed differs.

From the release's `lean` directory, the principal replay commands are

```
lake build
lake build CollatzEndpointTransport.Linear.PaperDependencyAudit
lake build CollatzEndpointTransport.Linear.PaperAudit
```

The first command reconstructs the retained source. The second reports transitive theorem dependencies. The third executes `#print axioms` on the paper-facing declarations; it reports only `propext`, `Classical.choice`, and `Quot.sound`. These checks supplement the written proof and are not used to discharge any step of it.

Acknowledgements

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 OpenAI ChatGPT and Codex for exploratory proof discussion, Lean assistance, literature-search assistance, and editorial review. Model outputs were treated as suggestions rather than mathematical evidence. The author checked every retained proof, computation, citation, and formal declaration and assumes full responsibility for the manuscript.

Funding

The author received no specific funding for this work.

Competing interests

The author declares no competing interests.

References

1. J. Allikvere, "Almost all Collatz orbits attain almost bounded values in natural density," preprint, version 2, 22 July 2026, doi:10.5281/zenodo.21499244.
2. J.-P. Allouche, "Sur la conjecture de Syracuse–Kakutani–Collatz," *Séminaire de Théorie des Nombres*, 1978–1979, Exp. No. 9, 15 pp., CNRS, Talence, 1979.
3. C. J. Everett, "Iteration of the number-theoretic function $f(2n) = n,\ f(2n+1) = 3n+2$," *Adv. Math.* 25 (1977), 42–45.
4. M. Inselmann, "An approximation of the Collatz map and a lower bound for the average total stopping time," arXiv:2402.03276v3.
5. I. Korec, "A density estimate for the $3x + 1$ problem," *Math. Slovaca* 44 (1994), 85–89.
6. L. Mazur, "Natural-density almost-bounded Collatz orbits in logarithmic time," preprint, version 2, 21 July 2026, https://www.proofatlas.ai/papers/natural-density-log-time-collatz/Mazur_Natural_Density_Collatz_Orbits_in_Logarithmic_Time_v2.pdf.
7. T. Tao, "Almost all orbits of the Collatz map attain almost bounded values," *Forum Math. Pi* 10 (2022), e12, doi:10.1017/fmp.2022.8.
8. R. Terras, "A stopping time problem on the positive integers," *Acta Arith.* 30 (1976), 241–252.