POSITIVE LOWER DENSITY OF COLLATZ PREDECESSORS

LECH MAZUR

Abstract. For the ordinary Collatz map, we prove that the set of positive integers whose orbit reaches a fixed positive target a has positive lower natural density whenever $3 \nmid a$. No assumption that a reaches 1, or that a is nonperiodic, is required. The proof combines weighted inverse-orbit estimates from earlier work with two arguments given here: every eligible target has a suitable nonperiodic predecessor, and inverse paths ending at that predecessor give distinct starting integers. The constants may depend on a. As a consequence, if any positive Collatz orbit fails to reach 1, the set of such starting values has positive lower natural density. Thus convergence to 1 on a set of upper natural density one would imply the Collatz conjecture. The main theorem and these implications have a Lean 4 formalization.

1. Introduction

Let C be the ordinary Collatz map on positive integers:

$$C(n) = \begin{cases} n/2, & 2 \mid n, \\ 3n+1, & 2 \nmid n. \end{cases} \tag{1.1}$$

Write $x \rightsquigarrow a$ when $C^t(x) = a$ for some integer $t \geq 0$, and put

$$\mathcal{P}(a) = \{x \geq 1 : x \rightsquigarrow a\}, \qquad A(X) = \#\{x \in A : 1 \leq x < X\} \quad (A \subseteq \mathbb{N}_{>0}).$$

All cutoffs X in density statements below are positive integers. We use $\underline{d}(A) = \liminf_{X\to\infty} A(X)/X$ and $\overline{d}(A) = \limsup_{X\to\infty} A(X)/X$. We write log for the natural logarithm. For a prime p and a positive integer n, $v_p(n)$ is the largest integer $e \geq 0$ such that p^e divides n. A unit residue modulo a power of 3 is a residue not divisible by 3.

Theorem 1.1. *For every positive integer a with $3 \nmid a$, there are a real number $c_a > 0$ and an integer X_a such that*

$$\mathcal{P}(a)(X) \geq c_a X \qquad \text{for every integer } X \geq X_a. \tag{1.2}$$

In particular, $\underline{d}(\mathcal{P}(a)) > 0$.

The target may be even. It need not be known to reach 1, and no classification of possible Collatz cycles is assumed. The conclusion is an eventual lower bound at *every* cutoff, not merely at a sequence of cutoffs. It neither asserts that the natural-density limit exists nor gives a common density bound for all targets.

The restriction $3 \nmid a$ is necessary. If $3 \mid a$, every predecessor step on a path to a must be a halving step: an odd step has output congruent to 1 modulo 3. Induction backwards along the path gives $\mathcal{P}(a) = \{2^j a : j \geq 0\}$, a set of density zero.

Let $\mathcal{G} = \mathcal{P}(1)$ and $\mathcal{E} = \mathbb{N}_{>0} \setminus \mathcal{G}$. Here failure to converge means failure to reach 1; it includes both a hypothetical nontrivial periodic orbit and an orbit escaping every finite set. The following consequences do not assert that either type exists.

Corollary 1.2. *Either every positive integer reaches 1, or $\underline{d}(\mathcal{E}) > 0$. Consequently, $\overline{d}(\mathcal{G}) = 1$ implies that every positive integer reaches 1.*

The second assertion is a conditional criterion, not a proof of its statistical hypothesis. Positive density of $\mathcal{G}$ alone, including the stronger time-controlled conclusion of the preceding paper, does not establish that hypothesis.

Date: September 6, 2026; version 1.0.
2020 *Mathematics Subject Classification.* Primary 11B83; Secondary 37P99, 03B35.
Key words and phrases. Collatz map, predecessor sets, lower natural density, orbit invariance, formal verification.

1.1. **Relation to the preceding paper.** Earlier work [2] proves positive lower natural density of starts reaching 1 within $10.46 \log n$ ordinary steps, with explicit density and cutoff constants. We use its mixing estimates and inverse-path construction. The new contribution is the transfer from one convergent seed to a seed reaching an *arbitrary* eligible target, together with the counting and exceptional-set consequences.

The preceding main theorem alone does not imply Theorem 1.1: positive density of starts reaching 1 supplies no path from those starts to an arbitrary a. Instead, we reuse the underlying estimates before convergence to 1 is imposed. Two new ingredients make this possible. First, every target $3 \nmid a$ has arbitrarily large odd predecessors in each prescribed residue class modulo a power of 3, and these predecessors can be chosen to have no positive return under the odd-only Syracuse map. Second, this no-return property supplies precisely the uniqueness needed in the count of distinct starting integers. We give these arguments below and state explicitly which analytic estimates are inherited.

This note establishes the qualitative fixed-target density statement and its exceptional-set consequences. General-target time bounds, including a $10.46 \log n$ bound, and explicit target-to-constants recipes are outside its scope.

1.2. **Context.** Krasikov and Lagarias [1, Theorem 6.1] obtained an eventual $X^{0.84}$ lower bound for each fixed positive target not divisible by 3, using the map that combines an odd step with one division by 2. The author's separate certified predecessor bounds [3] give eventual lower bounds $X^{0.90}$ and $c'_a X^{0.901}$ with $c'_a > 0$ depending on the target; they are context, not inputs to the present proof. Wirsching's preliminary manuscript [6] discusses positive predecessor density and a stronger uniform bound of the form c/a for noncyclic targets not divisible by 3. Theorem 1.1 addresses the qualitative fixed-target question; it does not establish that uniform bound or complete the particular proposed argument in [6].

The map convention matters for an individual even target. If $T(n) = n/2$ for even n and $T(n) = (3n+1)/2$ for odd n, and $\mathcal{P}_T$ denotes its predecessor set, then

$$\mathcal{P}(2a) \subseteq \mathcal{P}_T(a) \subseteq \mathcal{P}(a). \tag{1.3}$$

Indeed, a T orbit retains the terms of a C orbit except the intermediate even output immediately after each odd step. If an occurrence of $2a$ is retained, the next step reaches a; if it is omitted, that combined step ends at a. Thus the family of results in Theorem 1.1 also gives qualitative positive density for T, by applying it to $2a$. Same-target equality of the two predecessor sets is not being asserted.

Tao [5] proves that logarithmically almost every orbit attains values below any prescribed function tending to infinity. His Remark 1.4 also discusses a heuristic positive-density attraction to exceptional orbits. The growing-threshold theorem does not say that almost every orbit reaches the fixed value 1, and does not supply the hypothesis of Corollary 1.2. The present proof uses the finite ternary mixing framework developed in the companion paper, which builds on Syracuse-map methods; it does not use Tao's almost-boundedness conclusion as a substitute for fixed-target convergence. We make no claim of an exhaustive priority comparison.

1.3. **Proof outline.** The analytic estimates do not initially count integers. They give a positive weighted sum of inverse paths ending at a sufficiently large odd seed in one prescribed residue class. Section 2 shows that such a seed can be chosen to reach any eligible target, while never returning to itself. Section 3 states the exact weighted-path input used from the preceding paper; its derivation from that paper is collected in Appendix A.

The new counting argument is in Section 4. If two paths from the same starting integer reached the chosen seed at different times, that seed would return to itself. This rules out repeated representations of a starting integer. An elementary count in residue classes then removes the auxiliary residue weight from the analytic estimate and gives a linear lower bound at every large cutoff. Section 5 deduces the nonconvergence consequence and the conditional density-one criterion.

2. Nonperiodic predecessors in prescribed residues

For a positive odd integer r, write

$$S(r) = \frac{3r+1}{2^{v_2(3r+1)}}.$$

We say that r has *no positive return* if $S^t(r) \neq r$ for every $t \geq 1$. This excludes periodicity of r itself, not eventual entry into a periodic orbit. For example, $S(5) = 1$ and $S(1) = 1$: the seed 5 has no positive return, although its orbit enters a cycle immediately.

Lemma 2.1. *Let $f : \mathbb{N} \to \mathbb{N}$ and fix $u \in \mathbb{N}$. There is $B \in \mathbb{N}$ such that every $r > B$ whose f-orbit reaches u has no positive return under f.*

Proof. If u is not periodic, no periodic point can reach u. If u has a period $p > 0$, every periodic point whose orbit reaches u belongs to the finite cycle $\{u, f(u), \dots, f^{p-1}(u)\}$. To see the latter assertion directly, suppose $f^m(r) = u$ and $f^k(r) = r$ with $k > 0$. Then $f^{m(k-1)}(u) = r$. Reducing the exponent modulo p places r in the displayed finite set. Take B to be its maximum. □

Lemma 2.2. *If r has no positive return under a map f, and $f^d(x) = r = f^e(x)$, then $d = e$.*

Proof. If, for example, $d < e$, then $f^{e-d}(r) = r$, a contradiction. □

Lemma 2.3. *Let $a > 0$ with $3 \nmid a$. For every $q \geq 0$, every residue $y \in \mathbb{Z}/3^q\mathbb{Z}$, and every height L, there is a positive odd integer $R > L$ such that*

$$R \equiv y \pmod{3^q}, \qquad R \rightsquigarrow a, \qquad S^t(R) \neq R \quad (t \geq 1).$$

Proof. Choose $e = 2$ if $a \equiv 1 \pmod 3$ and $e = 1$ if $a \equiv 2 \pmod 3$, and put $r_0 = (2^e a - 1)/3$. For $k \geq 0$ define

$$R_k = \frac{(3r_0 + 1)4^k - 1}{3} = (3r_0 + 1)\frac{4^k - 1}{3} + r_0. \tag{2.1}$$

For $k \geq 1$ these are positive odd integers, and

$$3R_k + 1 = 2^{2k+e}a, \qquad C^{1+2k+e}(R_k) = a.$$

All R_k with $k \geq 1$ have the same next Syracuse value, namely the odd part $u = a/2^{v_2(a)}$ of a. In particular, this value is not assumed to equal a when a is even.

The elementary valuation identity $v_3(4^h - 1) = 1 + v_3(h)$ for $h > 0$ gives

$$v_3(R_{k+h} - R_k) = v_3(h).$$

For completeness, when $3 \nmid h$, expansion of $(1+3)^h$ modulo 9 gives valuation one. Replacing an exponent h by $3h$ increases the valuation by one, since for $z \equiv 1 \pmod 3$ the factor $z^2 + z + 1$ has valuation exactly one. This proves the identity. Thus $R_0, \dots, R_{3^q-1}$ represent all residues modulo 3^q exactly once, and adding 3^q to an index preserves its residue. The sequence tends to infinity. Apply Lemma 2.1 to S on the positive odd integers and the fixed endpoint u. A sufficiently large representative of y, beyond both L and the resulting bound, has no positive return. □

The lemma does not assert that these particular R_k already have positive density. Their role is to supply a single suitable seed for the inherited inverse-orbit estimates.

3. The weighted inverse-path input

We state the precise consequence of the preceding paper needed here. This separates the inherited analytic estimates from the new argument that turns weighted paths into a count of integers.

3.1. **Exact paths and their weights.** For a positive odd endpoint R, an *exact valuation word ending at R* is a finite sequence $w = (a_1, \dots, a_d)$ of positive integers for which some positive odd integer x_w satisfies

$$S^d(x_w) = R, \qquad a_i = v_2\big(3S^{i-1}(x_w) + 1\big) \quad (1 \leq i \leq d).$$

The entries are in *forward* order, from x_w toward R. The preceding paper lists the same entries in the reverse order for its inverse construction. Given R and w, the source x_w is unique: recover the orbit backwards using $x = (2^a z - 1)/3$ at each step. An exact word requires positive odd integer values throughout, not merely congruences.

Write $d(w) = d$, $A(w) = \sum_{i=1}^{d} a_i$, and

$$\omega(w) = 3^{d(w)}2^{-A(w)}. \tag{3.1}$$

The empty word, if present, has source R and weight 1. Composition of the affine maps $z \mapsto (3z+1)/2^{a_i}$ gives $2^{A(w)}R = 3^{d(w)}x_w + B_w$ with $B_w \geq 0$. In particular,

$$x_w\omega(w) \leq R. \tag{3.2}$$

Different words need not yet have different sources: an orbit could return to R. Section 4 will settle that issue.

3.2. **The analytic statement.** For $k \ge 1$, put $G_k = \mathbb{Z}/3^k\mathbb{Z}$ and $\langle f\rangle_k = 3^{-k}\sum_{z\in G_k} f(z)$. Let $\pi_{k,m} : G_k \to G_m$ be reduction when $1 \le m \le k$. Functions on G_k are also evaluated at integers by reduction modulo 3^k. Define the integer affine maps

$$\Psi_j(x) = 4^j x + \frac{4^j - 1}{3} \qquad (j \ge 0). \tag{3.3}$$

Proposition 3.1 (Input from the preceding paper)**.** *There exist positive integers B_0, C_*, q, a unit residue $y \in G_q$, and nonnegative functions $\rho_k : G_k \to \mathbb{R}$ for $k \ge 1$ with the following properties.*

First, ρ_k vanishes on nonunits, has mean $2/3$, and

$$0 \le \rho_m \le \frac{2}{3}\,3^m, \qquad \big\langle |\rho_k - \rho_m \circ \pi_{k,m}| \big\rangle_k \le \frac{C_*}{m^2} \quad (1 \le m \le k). \tag{3.4}$$

Second, for every odd $R \ge B_0$ with $R \equiv y \pmod{3^q}$ and every integer $m \ge 1$, there is a real $H = H(R,m) \ge 1$ such that every *real $X > H$ has the following data: integers $k \ge m$ and $J \ge 0$, a finite set $\mathcal{W}$ of exact valuation words ending at R, and selected weights v_w satisfying*

$$3^k \le X, \qquad X \le x_w < 32X, \qquad 0 \le v_w \le \omega(w) \quad (w \in \mathcal{W}), \tag{3.5}$$

for which

$$\frac{2}{3} \le \sum_{w\in\mathcal{W}} v_w \sum_{j=0}^{J} 4^{-j}\rho_k\big(\Psi_j(x_w)\big). \tag{3.6}$$

No convergence or no-return condition on R is assumed.

The functions ρ_k are auxiliary residue weights. The sum in (3.6) is not a count of integers, nor does its lower bound assert that the x_w are distinct. The selected weights retain the restrictions used in the preceding construction. Sources x_w may be divisible by 3, even though R is a unit and the functions ρ_k vanish on nonunits.

Proposition 3.1 collects finite-distribution mixing, residue selection, last-step reduction, and coverage of all large scales from [2]. Appendix A derives this statement with references to the estimates used and without using a distinct-source estimate. It is an expository combination of those inputs, not an additional independently formalized declaration.

4. Distinct sources and the general-target theorem

4.1. **A count in residue classes.** The next two elementary estimates provide the new counting step. They apply to any finite set of exact valuation words; no part of their source-uniqueness conclusion is assumed in Proposition 3.1.

Lemma 4.1 (Distinct starting integers)**.** *Let R have no positive Syracuse return, and let $\mathcal{W}$ be a finite set of exact valuation words ending at R. Then $w \mapsto x_w$ is injective. If $x_w \ge X > 0$ for every w, then for any nonnegative function g on these sources,*

$$\sum_{w\in\mathcal{W}} \omega(w) g(x_w) \le \frac{R}{X} \sum_{x\in\{x_w : w\in\mathcal{W}\}} g(x). \tag{4.1}$$

Proof. If $x_w = x_{w'}$, both lengths are hitting times of R from that source. Lemma 2.2 makes the lengths equal. The deterministic forward Syracuse orbit then fixes every valuation, so $w = w'$. Now use (3.2) to obtain $\omega(w) \le R/x_w \le R/X$, and sum over distinct sources. □

Lemma 4.2 (Residue-class estimate)**.** *Suppose the hypotheses of Lemma 4.1 hold, $x_w \in [X, 32X)$, $3^k \le X$, and $0 \le v_w \le \omega(w)$. For every nonnegative function $g : G_k \to \mathbb{R}$,*

$$\sum_{w\in\mathcal{W}} v_w g(x_w) \le 33R\langle g\rangle_k. \tag{4.2}$$

Proof. Each residue class modulo 3^k has at most $32X/3^k + 1 \le 33X/3^k$ nonnegative integers below $32X$. By Lemma 4.1, the left side is bounded by R/X times the sum of g over the distinct sources. Enlarge that nonnegative sum to all possible integers in $[X, 32X)$ and apply the residue-class bound. □

4.2. **Removing the auxiliary residue weight.** Fix a seed R satisfying Proposition 3.1 and having no positive return. Choose

$$m = 132RC_* + 1, \tag{4.3}$$

and for any $X > H(R, m)$ take its data from that proposition. Put

$$U = \sum_{w \in \mathcal{W}} v_w, \qquad \delta = |\rho_k - \rho_m \circ \pi_{k,m}|.$$

Since $\rho_m \le (2/3)3^m$ and $\sum_{j=0}^{J} 4^{-j} \le 4/3$, the contribution of the lifted ρ_m to (3.6) is at most $(8/9)3^m U$. The same selected weights v_w remain in U. Here mixing compares two periodic functions at different moduli; it does not make the residue weights constant.

For the error, each Ψ_j permutes the *full* group G_k: its slope 4^j is invertible modulo 3^k. Hence $\langle \delta \circ \Psi_j \rangle_k = \langle \delta \rangle_k$. Lemma 4.2 gives

$$\sum_{w \in \mathcal{W}} v_w \delta(\Psi_j(x_w)) \le 33R\langle \delta \rangle_k.$$

Multiply by 4^{-j}, sum, and use (3.4). Together with (3.6), this proves

$$\frac{2}{3} \le \frac{8}{9} 3^m U + \frac{44RC_*}{m^2}. \tag{4.4}$$

The choice (4.3) has $m^2 \ge 132RC_*$, so the last term is at most $1/3$. Therefore

$$U \ge \eta := \frac{3}{8\,3^m} > 0. \tag{4.5}$$

The loss estimate thus applies to the same selected paths that will be counted. Only the nonnegative error was enlarged to all possible sources in its residue classes.

4.3. **Every sufficiently large cutoff.**

Proof of Theorem 1.1. Fix $a > 0$ with $3 \nmid a$. Choose B_0, C_*, q, y from Proposition 3.1. By Lemma 2.3, there is one odd $R \ge B_0$ in that residue class which reaches a and has no positive Syracuse return. Choose m by (4.3) and let H be an integer at least $H(R, m)$. These choices all precede the cutoff: R and m will remain fixed.

For any integer $Y \ge 32(H + 1)$, set $X = Y/32$. The proposition supplies its selected words, whose sources lie in $[Y/32, Y)$ and reach R, hence a. Lemma 4.1, with $g = 1$ and $v_w \le \omega(w)$, and (4.5) give

$$\mathcal{P}(a)(Y) \ge \#\{x_w : w \in \mathcal{W}\} \ge \frac{XU}{R} \ge \frac{\eta}{32R} Y.$$

Thus Theorem 1.1 holds with

$$c_a = \frac{3}{256R\,3^m} > 0, \qquad X_a = 32(H + 1). \tag{4.6}$$

Only the finite path data—the word family, selected weights, k, and J—vary with Y; the seed, positive coefficient, and cutoff do not. □

Remark 4.3 (Scope of the constants). Lemma 2.1 selects a bound by distinguishing whether a fixed endpoint is periodic. It gives no effective bound on a possible cycle in this argument. The resulting choice of R is existential, so (4.3) and (4.6) do not by themselves give an effective target-to-constants algorithm or a bound uniform in a. General-target quantitative estimates and time bounds are outside the scope of this note.

5. The nonconvergence consequence

Every positive orbit reaches a positive integer prime to 3. Indeed, write $n = 2^j u$ with u odd. After j halvings and one odd step, the orbit reaches $3u + 1$, which is positive and not divisible by 3.

Reaching 1 is invariant in both directions along a finite path. Backward invariance follows by concatenating paths. For forward invariance, if a path has already passed through 1 before it reaches a, then a belongs to the known cycle $1 \to 4 \to 2 \to 1$ and still reaches 1. Otherwise the remaining path from a reaches 1. Hence failure to reach 1 is invariant in both directions as well.

Proof of Corollary 1.2. If $\mathcal{E}$ is nonempty, choose $n \in \mathcal{E}$ and then a positive target a with $n \rightsquigarrow a$ and $3 \nmid a$ as above. Invariance gives $a \in \mathcal{E}$ and $\mathcal{P}(a) \subseteq \mathcal{E}$. Theorem 1.1 therefore gives $c > 0, X_0$ such that $\mathcal{E}(X) \geq cX$ for every $X \geq X_0$. If $\overline{d}(\mathcal{G}) = 1$, there is a positive $X \geq X_0$ with $\mathcal{G}(X) \geq (1 - c/2)X$. Since the two sets are disjoint subsets of the integers below X, their counts sum to at most X. But the two lower bounds sum to $(1 + c/2)X > X$, a contradiction. □

Equivalently, the sufficient statistical hypothesis can be stated without density notation: for every $\varepsilon > 0$ and every integer N, there is an integer $X \geq N$, $X > 0$, such that $\mathcal{G}(X) \geq (1 - \varepsilon)X$. This is the exact form used in the formal theorem. If every positive integer reaches 1, the hypothesis plainly holds because $\mathcal{G}(X) = X - 1$. Thus it is also necessary. No finite computational check establishes its arbitrarily small error and arbitrarily large cutoff quantifiers.

Bidirectional invariance is essential to this argument. Properties measured relative to the initial value, such as descent below that value or failure of a particular logarithmic-time bound, are not automatically invariant. Conclusions about those properties cannot be inserted in place of convergence to the fixed target 1.

6. Formal verification and scope

The general-target theorem, the nonconvergence consequence, and the conditional convergence criterion have formal proofs in Lean 4.30.0-rc2, using the Lean mathematical library [4]. The main formal statement counts distinct positive integers $n < X$ reaching a, with exactly $a > 0$ and $3 \nmid a$ as its target hypotheses.

Recorded audits of the underlying development report only the standard Lean axioms of propositional extensionality, classical choice, and quotient soundness, or subsets thereof, and no admitted proof or unproved project axiom in the audited dependencies. Those checks reused compiled dependencies. The accompanying source supplement retains and adapts the portions needed for the three stated conclusions. Fresh compilation of its ten implementation/support modules and public interface against the slim version 2.1 supporting library, followed by an exact-statement and axiom audit, passed on September 7, 2026. All three public theorems reported only the three standard axioms listed above. This check reused compiled supporting dependencies; no full-library rebuild or independent LeanChecker run is claimed. The accompanying supplement *Verification and Reproduction* (`VERIFICATION.md`) gives the exact theorem correspondence, source provenance, verification evidence, and reproduction instructions. The paper and formalization page and source page provide stable links for this release.

Formal checking does not mechanically verify this exposition or its historical comparisons. Proposition 3.1 combines cited analytic statements; it is not advertised as a separately checked formal declaration. The map comparison (1.3) and the elementary converse of the statistical criterion are likewise proved in the text without being advertised as separately audited declarations.

Authorship and generative-AI disclosure. Lech Mazur directed the research and is responsible for the manuscript and its claims. The ProofAtlas.ai harness and OpenAI GPT models accessed through Codex contributed substantially to mathematical exploration, proof development, formalization, checking, and exposition. Anthropic Claude reviewed the preceding paper. The present manuscript underwent automated mathematical and source-correspondence reviews, not human peer review. AI output and agreement between models are not themselves proof evidence.

Appendix A. Derivation of the weighted-path input

This appendix explains exactly which parts of [2] give Proposition 3.1. The references below are to its published version 2.0. The argument stops before any estimate that requires distinct integer sources; that estimate was proved locally in Section 4.

A.1. **The reference functions and mixing.** The functions in [2, Section 2, equations (2.2)–(2.3)] are

$$\rho_k(z) = \frac{2}{3}\, 3^k \mu_k(z),$$

where μ_k is an auxiliary probability distribution on G_k. It is obtained by taking independent positive integers E_i with $\mathbb{P}(E_i = e) = 2^{-e}$ and iterating

$$Y_0 = 0, \qquad Y_{k+1} = 2^{-E_{k+1}}(3Y_k + 1) \pmod{3^{k+1}}.$$

Here Y_k is a residue modulo 3^k and $2^{-E_{k+1}}$ denotes the modular inverse of the indicated power of 2. The distribution of Y_k is μ_k.

Consequently $\rho_k \geq 0$, $\langle \rho_k \rangle_k = 2/3$, and ρ_k vanishes on nonunits for $k \geq 1$. Since a probability at one residue is at most one, $\rho_m \leq (2/3)3^m$. Equation (2.3) of [2], with exponent 2, gives the second inequality in (3.4) after enlarging the coefficient to a positive integer C_*. Its exponent-6 form is also used in the uniform estimates recalled next. This random construction defines the reference weights; it makes no independence assertion about valuations along a fixed integer orbit.

A.2. **A residue chosen before its representative.** Use the enlarged allowances from [2, Sections 8.2–8.3]:

$$b = 2^{80}, \qquad B_0 = 16^b, \qquad b_0 = b, \quad b_{n+1} = b_n + \lfloor b_n/100 \rfloor, \quad K_n = 16 + \lfloor n/100 \rfloor, \quad k_n = \lfloor b_n/4 \rfloor.$$

A construction stage consists of a fixed number of successive blocks of inverse steps. Each block ends at the prescribed first crossing of the barrier in [2, equation (3.1)]; its excess above that barrier is at most K_n. The central length restrictions from its Section 3 and the last-block restriction in its equation (5.2) are retained as indicator functions, not used to renormalize weights.

At a fixed stage N, pulling the reference weight backwards gives a nonnegative function g_N of the initial seed modulo a finite 3^q, with $q \geq 1$. Its mean over all residues is $(2/3)\prod_{j<N} p_j$, and it vanishes on nonunits. Here p_j is the geometric probability of an allowed central block. Equation (8.7) of [2] gives $\prod_{j<N} p_j \geq 255/256$. Units constitute exactly two thirds of the residues, so the mean of g_N over units is that product. Some unit y therefore has $g_N(y) \geq 255/256$.

The choice of N can be made before choosing an integer representative of y. The central variation estimate [2, Proposition 3.1] and the seed-independent bounds in its Section 8.3 control the later change for every sufficiently large odd representative. They bound the total central variation by $1/16$, the last-block comparison error by $1/8$, and the excluded last-block mass by $1/8$. Thus the residue-weighted selected sum at every sufficiently late stage satisfies

$$M_{n,X}(R) \geq 255/256 - 1/16 - 1/8 - 1/8 = 175/256 > 2/3. \tag{A.1}$$

This is the calculation in [2, equation (8.8)], used before specializing the seed to one that reaches 1. Its uniform thresholds depend on the initial scale and fixed mixing estimates, not on the chosen numerical representative R or the allowed external scale X.

For precision, $M_{n,X}(R)$ sums $\omega(h)\theta_c(h)\theta_t(h)\rho_{k_n}(x_h)$ over the stage's inverse paths h ending at R, including their prescribed block divisions. The functions θ_c, θ_t are respectively the central and last-block indicators just described. This sum uses the last-block excess allowance K_n. The shorter-excess paths used in Proposition 3.1 are obtained by the next reduction.

A.3. **Reduction to paths in a fixed size interval.** Use the reduction in the proof of [2, Proposition 5.3], before that proof invokes a source-counting estimate. Decrease the last valuation in *inverse* order by $2j$ until the excess above the crossing barrier is 0 or 1. The valuation remains positive, since the preceding partial sum was below the barrier. In our forward word convention this changes the *first* valuation. The earlier inverse blocks and the length of the last block are unchanged, so both selection indicators are retained.

If the reduced source is x and its weight is ω, the original source and weight are respectively $\Psi_j(x)$ and $4^{-j}\omega$, with $0 \leq j \leq J := \lfloor K_n/2 \rfloor$. The reduced paths have exact positive odd sources in $[X, 32X)$, by the endpoint bounds in [2, Lemma 6.1 and equation (5.1)]. Their sources need not be units. For this step only those endpoint bounds are used, not the source-uniqueness assertion of that lemma.

At one fixed stage, the seed and the whole inverse word recover all block divisions: each division is the specified first crossing. The final block is the remaining suffix after the preceding divisions. Thus replacing these divided paths by a finite set of their whole words neither loses a copy of a weight nor identifies two differently divided paths. Reversing the words gives the forward convention of Section 3. This recovery does not assume that a source reaches R at a unique time, and it makes no claim about comparisons between different stages.

Let $\mathcal{W}$ be this finite set of reduced words and put $v_w = \omega(w)\theta_c(w)\theta_t(w)$. The reduction is injective into the pairs consisting of a reduced word and its reconstruction index j. Enlarging its image to all such pairs is allowed by nonnegativity. Hence

$$M_{n,X}(R) \leq \sum_{w \in \mathcal{W}} v_w \sum_{j=0}^{J} 4^{-j} \rho_{k_n}(\Psi_j(x_w)).$$

Together with (A.1), this is (3.6). In particular, the original residue-weighted sum and the reduced-path sum have not been equated by simply deleting a residue factor.

A.4. **The modulus bound and all large scales.** After R and m are fixed, choose a stage threshold n_0 beyond the uniform threshold in (A.1), with $m \le k_n \le b_n$ for all $n \ge n_0$. Increase n_0 to satisfy the overlap threshold in [2, Lemma 6.2]. That lemma supplies one $H = H(R, m)$ such that every real $X > H$ belongs to the admissible scale interval of every endpoint at some stage $n \ge n_0$. This conclusion is for every sufficiently large scale, not merely for one subsequence.

Positivity in (A.1) gives a contributing path. The lower endpoint estimate in the proof of [2, Proposition 5.3], before any residue count, yields

$$X > \frac{16^{b_n}}{4 \cdot 3^{b_n}} \ge 3^{b_n} \ge 3^{k_n}.$$

Thus $k = k_n$ meets $3^k \le X$ as well as $k \ge m$. The chosen stage, reduced word family, selected weights, and $J = \lfloor K_n/2 \rfloor$ give all of Proposition 3.1. These steps have used neither convergence of the seed nor injectivity of the map from words to starting integers.

References

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

[2] L. Mazur, *Explicit Positive-Density Collatz Convergence in Logarithmic Time*, version 2.0, ProofAtlas, September 5, 2026, with accompanying Lean formalization. Available online: published paper and formalization.

[3] L. Mazur, *Certified* $x^{0.90}$ *Lower Bounds for Collatz Predecessor Sets*, version 2, ProofAtlas, July 17, 2026. Project release.

[4] The mathlib Community, *The Lean mathematical library*, Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (2020), 367–381. arXiv:1910.09336.

[5] 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.

[6] G. J. Wirsching, *On positive predecessor density in* $3n + 1$ *dynamics*, undated preliminary manuscript. Author's PDF.

Email address: 2022melodies@gmail.com