

Positive density of Collatz convergence at every rate above $3/\log(4/3)$

Idris Ali Shaik

Independent researcher (unaffiliated)

shaikidris@apache.org

ORCID: 0009-0009-9699-9712

Preprint, version 1.0.0, 21 September 2026

[doi:10.5281/zenodo.22871622](https://doi.org/10.5281/zenodo.22871622)

© 2026 Idris Ali Shaik. Licensed under CC BY 4.0.

Abstract

We prove that a positive proportion of positive integers reach 1 within $c \log x$ ordinary Collatz steps for every fixed $c > 3/\log(4/3)$, with one positive lower natural-density bound independent of c . For each fixed positive target not divisible by three, the same range of times holds with a target-dependent bound, irrespective of its subsequent orbit. Targets divisible by three have predecessor basins of natural density zero. Diagonalization gives one positive-density set with a vanishing excess above the limiting time coefficient.

The proof extends weighted inverse trajectories from a fixed finite family of roots and preserves positive weighted mass at every sufficiently large scale. We derive an explicit Fourier-decay estimate, at frequencies not divisible by three, for an auxiliary distribution on residue classes. Its proof uses paired geometric valuations and their exact binomial first-passage law. At order 6409, the coefficient C satisfies $\log_2 \log_2 \log_2 C < 176$, compared with $> 2^{172}$ for the explicit coefficient in Mazur's work. Fractional-moment certificates, residue-capacity bounds, and an enlarged family of roots give explicit density estimates after harmonic and dyadic counting.

The primitive-decay theorem and the main qualitative density conclusions are formalized in Lean 4. The coefficient-size comparison and the complete optimized density bounds remain paper proofs.

Keywords: Collatz problem; hitting times; lower natural density; inverse trajectories; Fourier decay; formal verification.

2020 Mathematics Subject Classification: 11B83, 37P99.

Introduction and main results

Ordinary hitting times and lower density

Write $\mathbb{N} = \{1, 2, 3, \dots\}$. The ordinary Collatz map is

$$\text{Col}(x) = \begin{cases} x/2, & x \text{ even,} \\ 3x + 1, & x \text{ odd.} \end{cases}$$

The Collatz conjecture asserts that every positive integer eventually reaches 1. Throughout this paper, convergence means reaching 1, and each application of Col counts as one step.

For a positive target y , let $\tau_y(x)$ be the least nonnegative integer j such that $\text{Col}^j(x) = y$, with $\tau_y(x) = \infty$ when there is no such j . All unmarked logarithms are natural, and $\log_2$ denotes base two. Put

$$c_0 = \frac{3}{\log(4/3)}, \quad \mathcal{G}_c(y) = \{x \geq 1 : \tau_y(x) \leq c \log x\}, \quad \mathcal{G}_c = \mathcal{G}_c(1).$$

The notation

$$\mathcal{G} = \{x \geq 1 : \tau_1(x) \leq (10431/1000) \log x\}$$

is used for the fixed numerical specialization. For $S \subseteq \mathbb{N}$, its lower natural density is

$$\underline{d}(S) = \liminf_{X \rightarrow \infty} \frac{\#(S \cap [1, X])}{X}.$$

Our results count both even and odd integers and give lower bounds at every sufficiently large cutoff. A positive lower natural density is weaker than density one and does not assert that the natural density exists. For a starting integer greater than one, reaching one also witnesses descent; our conclusions control the time to reach one, rather than only a drop below the starting value.

Main results

Theorem I.1 (positive density at every rate above c_0). There is an explicitly computable constant $c_* > 0$ such that, for every real $c > c_0$,

$$\liminf_{Y \rightarrow \infty} \frac{1}{Y} \#\{x \in \mathbb{N} : x \leq Y, \tau_1(x) \leq c \log x\} \geq c_* > 0. \quad (\text{I.main})$$

The constant c_* is independent of c . For each fixed $c > c_0$ and $0 < d < c_*$, the proportion is at least d at every sufficiently large cutoff; that cutoff may depend on c and d . The finite certificates and effective finite-search approximation defining c_* are given in (D.densityrecipe).

The ratio $c_0 \approx 10.42817849$ arises from geometric valuations with $\Pr(K = k) = 2^{-k}$, $k \geq 1$. One accelerated odd step has mean ordinary cost $1 + \mathbb{E}K = 3$ and mean logarithmic contraction $\mathbb{E}K \log 2 - \log 3 = \log(4/3)$. Related stochastic stopping-time models are surveyed in [5]. The exact stopped means (S.means) and the weighted inverse construction connect this ratio to the counting theorem. The **endpoint discussion** describes the role of the strict inequality $c > c_0$.

Theorem I.2 (every fixed target coprime to three). For every positive integer y with $3 \nmid y$, there is an effectively defined $d_y > 0$ such that

$$\forall c > c_0 : \quad \underline{d}(\mathcal{G}_c(y)) \geq d_y. \quad (\text{I.targets})$$

The estimate applies irrespective of the orbit after the target is reached. Its constant is fixed before c . The **fixed-target construction** gives the root-score recipe and its one-period refinement (D.oneperiodscore). The restriction on y is exact: when $3 \mid y$, its predecessor basin is $\{2^j y : j \geq 0\}$, which has natural density zero. Disjoint predecessor basins also preclude a positive constant uniform over all admissible targets.

An elementary maximum estimate gives the common-clock result without the large finite moment tables:

$$\forall c > c_0 : \quad \underline{d}(\mathcal{G}_c) \geq d_{\text{qual}} > 0. \quad (\text{I.qualitative})$$

Both the root score and the coarse residue level in (D.elementarydensity) are chosen before c . The qualitative argument can also use Tao's reference-law mixing estimates [1]. The common density has the diagonal consequence (D.vanishingclock): there is one positive-density set on which $\tau_1(x) \leq (c_0 + \eta(x)) \log x$ for a nonincreasing $\eta(x) \rightarrow 0$.

Root scores and explicit density bounds

The construction begins with a finite family of odd *roots* with specified paths to the target. A reference distribution modulo powers of three assigns a marker to each inverse trajectory. A *root score* $W > 0$ is a guaranteed sum of persistent marked weights, each divided by its root. Sections 7–8 construct the baseline rational score $W_* \in (0, 2^{-4096}]$.

There are four density estimates for a given score. The formula $c_{\text{MV}}(W)$ uses maximum and centered second-moment bounds; $c_{\text{FM}}(W)$ adds a direct 3/2-moment with twelve-step propagation; $c_{\text{FM},16}(W)$ also uses sixteen-step propagation; and $c_{\text{comp}}(W)$ uses complementary residue capacity. Section 10 defines these functions by optimization over the coarse residue level, with a common mixing envelope and counting conversion.

For every $0 < W \leq 2^{-4096}$, writing $a = \log_2(1/W)$, we prove

$$\frac{c_{\text{FM}}(W)}{c_{\text{MV}}(W)} > 2^{2^{(a-9)/95}/1000-2a-31} > 2. \quad (\text{I.gain})$$

At the same score, (Q.sixteengain) gives $c_{\text{comp}}(W) > c_{\text{FM},16}(W) > (1369/625)c_{\text{FM}}(W)$. These comparisons concern the complete optimized lower-bound formulas. Their absolute size is estimated in Section 11.

Quantitative refinement of the root pool. The residue-completion construction in **Appendix C** supplies a score $W_{\mathcal{F}}^{\text{joint}} = W_{\mathcal{F}}^{\text{joint}}(q_N, P)$, where q_N and P are the startup level and residual mass from (R.residual)–(R.score). It satisfies

$$W_{\mathcal{F}}^{\text{joint}} > \mathfrak{A}_b W_*, \quad \mathfrak{A}_b = 2.9273409933 \dots \times 10^{275}, \quad b = 2^{25}. \quad (\text{I.rootgain})$$

For comparisons within the twelve-step formula, define

$$\begin{aligned} c_* &:= c_{\text{FM}}(W_*), \\ c_{\text{amp}} &:= c_{\text{FM}}(\mathfrak{A}_b W_*), \\ c_{\text{joint}} &:= c_{\text{FM}}(W_{\mathcal{F}}^{\text{joint}}). \end{aligned} \quad (\text{I.constants})$$

The combined density bound is

$$\begin{aligned} \underline{d}(\mathcal{G}_c) &\geq c_{\text{comp}}(W_{\mathcal{F}}^{\text{joint}}) \\ &> c_{\text{joint}} > c_{\text{amp}} > c_* > 0 \quad (c > c_0). \end{aligned} \tag{I.rootdensity}$$

The first comparison follows from (Q.sixteengain), and the remaining ones from (Q.cubicscore) and (C.domainguard). All constants are independent of c . The root construction redistributes valuations along fixed trajectories and subtracts their sum from the incoming inverse step. This preserves the total ordinary clock for each seed and index. Backward residue completion and prefix-height estimates produce residue-complete nonrecurrent root blocks.

The counting mechanism

Exact affine formulas identify the valuation words that give positive integer trajectories. After normalization by the roots and identification of repeated representations, the aggregate weight of a source x is at most $1/x$. Thus source weight u in $[X, \Lambda X)$ already requires at least uX distinct integers. Harmonic allocation and counting their disjoint dyadic towers give the sharper conversion used in Section 10.

The analytic step bounds primitive Fourier coefficients of the reference law. Appendix E derives an explicit coefficient using the original pair process, its binomial first-passage law, and separated phase triangles. Section 6 converts this estimate into all-level mixing with exponent $2314/25$ and coefficient $C_6 = 2C_{\text{primitive}}20^{6409} + 2$. This controls the total variation of the root marks and the error in passing to a fixed coarse level.

Splitting a reference word gives a mixture of affine permutations. Jensen's inequality therefore bounds a moment at a combined level by the product of the moments at the two component levels, for every real order $p \geq 1$. The finite direct $3/2$ -moment certificates extend to all levels through this product inequality. Hölder's inequality, applied to both occupied and complementary residue capacity, converts marked mass into source mass.

For example, the second-moment bound $M_2(12) \leq 6.803$ and mean one give $M_{3/2}(12) \leq \sqrt{6.803} \approx 2.608256$, whereas direct calculation gives $M_{3/2}(12) \leq 2.113605$. The counting denominators from the second-moment block at sixteen and fractional block at twelve have rates

$$8.685^{1/16} \approx 1.14465106, \quad 2.113605^{1/6} \approx 1.13284538.$$

The slower growth of the latter compensates for its extra power of the residual marked mass, as (I.gain) shows. The direct depth-sixteen fractional certificate further reduces its denominator rate to $1.10650942918 \dots$

The seed words have exact depth b_j and total valuation in $[2b_j - \lceil b_j^{3/5} \rceil, 2b_j + \lceil b_j^{3/5} \rceil]$. Their reference survival is at least $24/25$. A summable variation bound then supplies a fixed positive root score. Varying the continuation and terminal parameters gives every fixed $c > c_0$ with this same score; $10431/1000 = 10.431$ is one numerical specialization.

Relation to previous work

Classical stopping-time and density results include the work of Terras [9], Everett [10], and Korec [11]. These concern descent along an orbit, whereas reaching a prescribed fixed target is the stronger event considered here. Lagarias [12] surveys the problem and its generalizations; Inslemann [13] studies approximation of the map and average total stopping time. These provide context for the distinction between a descent witness, a hitting-time bound, and a counting theorem for convergence.

Krasikov and Lagarias [6] use difference inequalities and a computer-assisted argument to obtain at least $X^{0.84}$ predecessors below X for each fixed positive target coprime to three, once X is sufficiently large. Tao [1] proves that the minimum of a Collatz orbit is below any prescribed function tending to infinity for almost all starting integers in logarithmic density. His proof develops the reference-law and Fourier methods used here.

Mazur [3] obtains convergence to 1 within $(523/50) \log x$ ordinary steps on a set of positive lower natural density, with an explicit finite cutoff. We build on his weighted inverse-trajectory and terminal-counting framework. The analytic and trajectory estimates used in our density proof are derived in this paper, including the primitive Fourier bound in Appendix E. The pairing and separated-triangle method follows Tao [1, Section 7] and its explicit development in [3]; our first-passage calculation works directly with the original pair process and supplies the coefficient used throughout.

The main clock result gives one density constant for every fixed $c > 3/\log(4/3)$. The moment estimates, complementary capacity, and root completion provide its quantitative refinements. Section 11 compares their formulas with the bound in [3], including its distinct finite-cutoff recipe. Appendix A compares the two seed constructions at a common analytic coefficient.

Mazur [4] also proves positive lower natural density of predecessors for each fixed positive target coprime to three. Theorem I.2 supplies ordinary hitting times throughout $c > c_0$, with one effective target-dependent constant. Its root construction applies independently of the orbit after the target is reached.

The author’s earlier preprints treat two related descent questions. The polylogarithmic-descent paper [7] gives natural-density-one conclusions with a common landing for the shortcut and ordinary orbits, logarithmic passage times, and control of the preceding trajectories. The endpoint-transport paper [8] proves an exceptional-set pullback estimate and applies it to stretched-logarithmic descent. Their target scales still grow with the starting value, so those results do not by themselves imply convergence to one. Here the target is fixed and the conclusion is positive lower natural density, rather than density one. The earlier papers are comparisons, not inputs to the proof below.

Terminology and recurring notation

A *source* is a starting integer being counted, a *target* is the integer to be reached, and a *root* is a selected odd intermediate integer with a specified finite path to that target. A *valuation word* records the positive numbers of halvings between successive odd states. The finite initial word families form the *seed*. A *graft* is the finite transition to the stopped blocks used to continue the inverse trajectories.


A *persistent mark* is a lower bound valid at every later stage. Its weighted sum, rather than each individual mark, is positive. A *residue-complete block* has one root in every unit residue class at the chosen modulus; *unit* means coprime to three. A root family is *nonrecurrent* when no positive accelerated iterate of a member belongs to the family, including the same member.

Notation	Meaning
$\text{Col}, \tau_y(x)$	Ordinary Collatz map and first hitting time; an odd step and every halving are counted separately.
c, c_0	Ordinary-step coefficient and the ratio $3/\log(4/3)$.
$c_*, c_{\text{amp}}, c_{\text{joint}}$	Density constants in (I.constants).
$G_q, \mathcal{U}_q$	All residues and unit residues modulo 3^q .
f_q, ρ_q, H_q	Mean-one reference density, its rescaling $(2/3)f_q$, and the reference fan.
W, W_*	A root score and the baseline score.
$c_{\text{MV}}, c_{\text{FM}}, c_{\text{FM},16}, c_{\text{comp}}$	Maximum/variance, twelve-step fractional, sixteen-step fractional, and complementary-capacity density formulas.
$m, k(X)$	Fixed coarse residue level and finer level growing with the cutoff.

Symbols confined to a construction have local meanings. The schedule coefficient L in Sections 3–5 is distinct from $L = 2 \log 2$ in counting; the seed sizes b_j are distinct from the growing terminal size b ; and the accelerated map U in Section 1 is distinct from the source mass U in Section 10. Local notation tables accompany the terminal and counting arguments.

Proof architecture and order of choices

The argument passes from analytic mixing to persistent root scores, then to integer source families and density:



The **quantitative reduction** states the source-family theorem and counting estimates first. Sections 1–2 establish the exact word identities and transfer bounds. Sections 3–5 construct all-scale sources from mixing and persistent-root estimates, proved in Sections 6–8. Sections 9–10 give the moment and counting arguments, and Section 11 compares their explicit bounds. The forward references from the construction are thus discharged before its final application.

The principal dependencies are summarized below. Equations retain their mnemonic labels; theorem numbers follow the section numbers, with I and R for the introduction and the preliminary reduction.

Conclusion	Inputs used	Proof location
Primitive decay	Exact pair law, first passage, phase geometry	Appendix E
Reference mixing	Primitive decay and finite Fourier decomposition	Section 6
Persistent positive root score	Mixing, seed survival, finite root period	Sections 7–8
Sources at every large scale	Exact word laws, persistent score, finite graft	Sections 1–5
Common-clock convergence and fixed-target hitting	Source families, capacity bounds, harmonic and dyadic counting	Section 10
Quantitative improvements	Direct moment certificates, coefficient comparison, larger root scores	Sections 9–11; Appendices B–C
Profile and entropy consequences	Affine mixtures, projectivity, summable mixing	Appendix D

Appendix A compares fixed-depth and first-crossing seeds. Appendix B proves monotonicity in the analytic coefficient. **Appendix C** establishes the root refinement; **Appendix D** develops concentration-profile and entropy consequences. **Appendix E** contains the primitive-decay proof required for the explicit theorem.

The analytic coefficient, seed, finite root pool, and moment certificates are fixed first. The root score and mixing envelope determine the coarse counting level, independently of the clock. Next choose $c > c_0$, its positive terminal parameter, an error tolerance, and a finite graft. The counting cutoff then tends to infinity. The root pool and graft remain fixed while the terminal trajectories and their precision grow. The final counting argument uses finitely many disjoint intervals before taking its remaining limits.

From source families to density

This section isolates the estimates used in counting. The intervening construction will supply the source families in the theorem below; its proof can be read after the quantitative purpose of those families is clear.

The reference distribution and source-family theorem

For an integer $m \geq 0$, let $G_m = \mathbb{Z}/3^m\mathbb{Z}$, and write $\langle h \rangle_m = 3^{-m} \sum_{r \in G_m} h(r)$. Take independent positive integers K_i with $\Pr(K_i = a) = 2^{-a}$, and put $S_i = K_1 + \dots + K_i$. The reference law μ_m is the distribution

of $\sum_{i=1}^m 3^{i-1}2^{-S_i}$ in G_m , where powers of two are inverted modulo 3^m . Its density relative to the uniform measure is $f_m(r) = 3^m \mu_m(r)$, and $\rho_m = (2/3)f_m$. The empty sum gives $f_0 = 1$. Section 2 relates this reference law to weighted integer trajectories by exact affine transfer identities.

The weighted sum needed after reducing terminal trajectories is

$$F_j(x) = 4^j x + (4^j - 1)/3, \quad H_m(x) = \sum_{j \geq 0} 4^{-j} \rho_m(F_j(x)).$$

We call this sum the reference fan. Each F_j permutes G_m , so $\langle H_m \rangle_m = 8/9 =: \mu$. The [reference-moment section](#) proves its exact relation to f_{m+1} and its moment estimates.

Theorem R.1 (all-scale source families). The root construction supplies a fixed rational score $W_* > 0$. For every fixed $c > c_0$, every $\eta > 0$ and every fixed $\Lambda > 16$, a single finite extension of the construction works at all sufficiently large real X . It gives a finite odd-source set $S_X \subseteq \mathcal{G}_c \cap [X, R_X X]$, weights $0 \leq a_x \leq 1/x$, and an integer $k(X)$, such that

$$W_* - \eta \leq \sum_{x \in S_X} a_x H_{k(X)}(x), \quad R_X < \Lambda, \quad R_X \rightarrow 16, \quad k(X) \rightarrow \infty, \quad 3^{k(X)}/X \rightarrow 0. \quad (\text{I.sources})$$

The roots and score are independent of c, η, Λ, X .

This is (N.everyclock), with the actual radius supplied by (N.selector) and (N.scales); (R.graftinput) proves its root and mixing premises. The score is defined by the finite recipe (R.score).

General moments give lower bounds for physical source mass

Fix a coarse level $m \geq 2$. Let ϵ_m be the mixing-error envelope (M.envelope), with its explicitly specified positive integer coefficient C_6 . In particular, for every $n \geq m$,

$$\langle |f_n - f_m \circ \pi_{n,m}| \rangle_n \leq \epsilon_m \leq C_6 m^{-2314/25},$$

where $\pi_{n,m}$ is reduction from G_n to G_m . Put $L = 2 \log 2$, $\kappa = (8/9)L$, and $p_m = (W - \kappa \epsilon_m)_+$, where $t_+ = \max\{t, 0\}$ and W is a score for which the source-family conclusion holds.

For $s > 1$, suppose $B_{s,m} \geq \langle H_m^s \rangle_m$. The total source mass $U = \sum_x a_x$ then has the limiting lower bound

$$v_{s,m}(W) = \frac{p_m^{s/(s-1)}}{(LB_{s,m})^{1/(s-1)}}. \quad (\text{I.generalp})$$

More precisely, every strict smaller nonnegative mass is available at every sufficiently large scale, after fixing the error allowance.

Here is the key inequality. With $Q = 3^m$, the actual residue masses $w_r = \sum_{x \equiv r \pmod{Q}} a_x$ obey

$$w_r \leq \frac{T_X}{Q}, \quad T_X = \frac{Q}{X} + \frac{1}{2} \log R_X \rightarrow L.$$

Indeed, odd representatives in a residue class have spacing $2Q$. Weighted Hölder therefore gives

$$\sum_r w_r H_m(r) \leq U^{1-1/s} \left(\sum_r w_r H_m(r)^s \right)^{1/s} \leq U^{1-1/s} (T_X B_{s,m})^{1/s}.$$

The fine-to-coarse error tends to at most $\kappa \epsilon_m$. Taking limits with m, s fixed yields (I.generalp). The proof of (D.generalp) below supplies the precise tolerance and all-scale quantifiers. The finite-cutoff estimates use T_X ; their limiting coefficient is L .

Distinct odd sources have disjoint sets of multiples by powers of two, and these multiples remain in $\mathcal{G}_c$, since $c \log 2 > 1$. Keeping their weights while counting gives the conversion

$$\mathcal{D}(u) = \frac{32}{225} (e^{2u} - 1), \quad 0 \leq u \leq \frac{1}{2} \log(3/2), \quad \underline{d}(\mathcal{G}_c) \geq \mathcal{D}(u). \quad (\text{I.conversion})$$

Here u must be available in every sufficiently late interval of every fixed ratio greater than sixteen. The proof is (D.radial). For the constructed range $0 < W \leq 27/2^{27}$, all mass bounds used here lie within this domain.

At $s = 3/2$, the direct certificate gives $B_{3/2,m} = \mathcal{B}_m$, where $\mathcal{B}_m^2 \leq 5(57/50)^m$, and (I.generalp) is $p_m^3/(L^2 \mathcal{B}_m^2)$. The formula $c_{\text{FM}}(W)$ retains the maximum of this branch and the maximum/variance branches, then optimizes the coarse level using the same conversion $\mathcal{D}$. Thus each of the long construction's inputs has a specific role: the root score supplies marked mass, mixing bounds the error in changing level, and concentration bounds force enough actual source mass to be counted.

Corollary R.2 (a simple rational density bound). For a rational source-family score $0 < W \leq 27/2^{27}$, define

$$m_o = \max \left\{ 2, \left\lceil \left(\frac{5C_6}{2W} \right)^{25/2314} \right\rceil \right\}.$$

Then

$$\underline{d}(\mathcal{G}_c) \geq c_{\text{FM}}(W) > \frac{4}{1125} W^3 \left(\frac{57}{50} \right)^{-m_o}. \quad (\text{I.rational})$$

The choice gives $p_{m_o} > W/2$; (D.explicit) proves the remaining constants and the integer-power recipe for m_o . This convenient rational bound is weaker than the optimized formula. In particular it applies to $W = W_*$, giving an explicit lower bound without first carrying out the full optimization.

1. Exact valuation words and physical trajectories

We first relate a finite list of valuations to a congruence class of trajectories and an exact affine formula. Reversing that formula constructs predecessors of a known convergent integer. The resulting weight bound, after histories with the same starting integer are combined, will be the input to the final count.

For a prime p and a positive integer n , write $\nu_p(n)$ for the largest integer $a \geq 0$ such that p^a divides n .

Let $U(x) = (3x+1)/2^{\nu_2(3x+1)}$ on positive odd integers. An ordinary Collatz step sends odd x to $3x+1$, and even x to $x/2$. One U -step of valuation k comprises $k+1$ ordinary steps.

Exact local certificates

For a word $w = (k_1, \dots, k_d)$ of positive integers, put

$$S_j = \sum_{i \leq j} k_i, \quad S_0 = 0, \quad B_0 = 0, \quad B_j = 3B_{j-1} + 2^{S_{j-1}}. \quad (\text{W.data})$$

A certificate $\text{WordCert}(r, E, w)$ consists of a positive odd representative r , an integer $E \geq S_d + 1$, and the finite checks $\nu_2(3U^{j-1}(r) + 1) = k_j$ for $1 \leq j \leq d$. An empty word is allowed, with $E \geq 1$.

Here a *source* is the starting integer of the forward trajectory, and its *endpoint* is the value reached after the word. The congruence condition $x \equiv r \pmod{2^E}$ is called a *cylinder*. Specifying a value to e bits of precision means fixing its residue modulo 2^e .

Theorem 1.1 (word invariance, affine endpoint, and precision). Every positive $x \equiv r \pmod{2^E}$ has exactly the certified word and satisfies

$$U^j(x) = \frac{3^j x + B_j}{2^{S_j}} \quad (0 \leq j \leq d). \quad (\text{W.affine})$$

For $x = r + 2^E h > 0$,

$$U^d(x) = U^d(r) + 3^d 2^{E-S_d} h. \quad (\text{W.output})$$

The guaranteed output precision is $E - S_d$ bits, and cannot uniformly be increased by one bit on this entire cylinder. If $2^{S_d} > 3^d$, descent holds for every such source with $x > B_d/(2^{S_d} - 3^d)$. The word takes exactly $S_d + d$ ordinary steps.

Proof. Suppose the two trajectories agree in valuations through step $j-1$. Their next pre-division numerators differ by $3^j(x-r)/2^{S_{j-1}}$, divisible by $2^{E-S_{j-1}}$. Since $E - S_{j-1} \geq k_j + 1$, this does not change the known valuation k_j . Division gives the next affine formula and difference, completing induction. The coefficient of h in (W.output) has exact valuation $E - S_d$, since 3^d is odd. Two consecutive sufficiently large nonnegative h 's give positive sources and endpoints differing by that valuation, so no stronger uniform precision is possible. Descent is exactly $(2^{S_d} - 3^d)x > B_d$. Summing $k_j + 1$ gives the clock. $\square$

Lemma 1.2 (one source cylinder per word). Every positive valuation word is realized by exactly one odd residue class modulo 2^{S_d+1} . That cylinder has Haar probability 2^{-S_d} among odd 2-adic integers.

In finite terms, exactly one of the 2^{S_d} odd residue classes modulo 2^{S_d+1} realizes the word. Haar probability assigns equal mass to these classes.

Proof. The empty word is the odd class modulo two. Suppose a prefix of total valuation S has its unique class modulo 2^{S+1} . By (W.output), refining it modulo 2^{S+k+1} makes its current endpoint run bijectively over odd residues modulo 2^{k+1} , because the variable coefficient is twice a unit. Exactly one residue obeys $3y + 1 \equiv 2^k \pmod{2^{k+1}}$, the condition for valuation k . That residue is odd, and the selected source class has positive representatives. Induction proves existence and uniqueness. The product of the successive Haar fractions 2^{-k} is 2^{-S_d} . The later counting argument uses the actual weights of integer sources. $\square$

Orientation and affine composition

An inverse word $(k_1, \dots, k_d)$ attached to an odd root M means

$$x_0 = M, \quad x_i = (2^{k_i} x_{i-1} - 1)/3. \quad (\text{W.inverse})$$

It is physically realized when every x_i is positive, odd, and integral. We use *physical* only to distinguish these actual integer trajectories from the auxiliary residue-class calculations introduced later. Then $U(x_i) = x_{i-1}$, and the chronological forward word from x_d is $(k_d, \dots, k_1)$, not the displayed inverse order. Direct substitution gives

$$M = 3^d 2^{-S_d} x_d + C_w, \quad C_w = \sum_{i=1}^d 3^{i-1} 2^{-S_i}. \quad (\text{W.offset})$$

The totals S_d, d and the clock $S_d + d$ are reversal invariant; prefix tests and additive offset formulas are not.

Proof. Solve the next inverse step as $x_{i-1} = 3 \cdot 2^{-k_i} x_i + 2^{-k_i}$. Substitution into $M = 3^{i-1} 2^{-S_{i-1}} x_{i-1} + C_{i-1}$ adds $3^{i-1} 2^{-S_i}$ to the offset and gives the next slope. Induction proves (W.offset). Oddness of x_{i-1} in $3x_i + 1 = 2^{k_i} x_{i-1}$ proves the valuation is exactly k_i . $\square$

If $M = \omega_1 y + C_1$ and $y = \omega_2 x + C_2$, then

$$M = \omega_1 \omega_2 x + C_1 + \omega_1 C_2. \quad (\text{W.compose})$$

This composition formula will be used when new inverse blocks are attached to a finite history. It retains the order of the blocks and their offsets.

Convergent roots and uniqueness of source weights

The *pre-1 odd basin* of a root consists of the odd integers whose forward trajectory visits that root before reaching one. A family of words is *prefix-free* if no selected word is a proper initial segment of another. *Parsing* a concatenation means identifying the cuts between its component words. These two properties will prevent a trajectory from being counted more than once.

For $s \geq 1$, let $R_s = (4^s - 1)/3$. It is positive and odd, $3R_s + 1 = 2^{2s}$, and $U(R_s) = 1$. For $s > 1$, its first ordinary hitting time of one is $2s + 1$; $R_1 = 1$ has hitting time zero. Distinct $R_s > 1$ have disjoint pre-1 odd basins: after visiting one of them a trajectory has next odd value one, where it remains. It cannot visit another root, or visit the same root twice before reaching one.

A physically realized complete history ending at $R_s > 1$, with source x and physical weight $\omega = 3^d 2^{-S_d}$, satisfies

$$x\omega \leq R_s \quad (\text{W.charge})$$

by the nonnegative offset in (W.offset). Under a deterministic prefix-free stopping/parsing rule, a source and root admit at most one complete retained history: the orbit and valuations are unique, the root is visited at most once, and parsing has no choice. Thus the same weight bound applies after grouping by actual source, provided compression preimages have first been identified rather than counted as additional reduced histories.

Normalize the weight at root R_s by dividing by R_s , and pool distinct such roots. Basin disjointness gives the aggregate weight bound $a_x \leq 1/x$, not one copy of that bound per root. The later terminal-compression proof must establish unique reduced-history parsing before invoking this conclusion.

Lemma 1.3 (source weights for a nonrecurrent root family). Let $\mathcal{P}$ be a finite set of distinct positive odd roots such that

$$U^j(M) \neq L \quad (M, L \in \mathcal{P}, j \geq 1), \quad (\text{W.antichain})$$

including $M = L$. Then a source visits the pool at most once. Under the same deterministic parsing and identification of compression preimages,

$$a_x = \sum_{M \in \mathcal{P}} \frac{1}{M} \sum_{\substack{z \text{ reduced from } M \\ x_z = x}} \omega_z \leq \frac{1}{x}. \quad (\text{W.poolcharge})$$

Indeed, two visits in chronological order would violate (W.antichain). The single visit fixes its valuation word and its parsed history, and $x\omega_z \leq M$ follows from (W.offset). Distinct convergent roots form such a pool whenever none is a forward iterate of another; a root greater than one that reaches one cannot recur. The nonrecurrence condition is imposed on the actual orbit. $\square$

Lemma 1.4 (one-step roots of an arbitrary odd unit). For $y > 0$ odd, $3 \nmid y$, define

$$e(y) = \begin{cases} 2, & y \equiv 1 \pmod{3}, \\ 1, & y \equiv 2 \pmod{3}, \end{cases} \quad X_t(y) = \frac{2^{2t+e(y)}y - 1}{3} \quad (t \geq 0). \quad (\text{W.constructor})$$

These are precisely the positive odd predecessors of y , without repetition, with exact valuation $2t + e(y)$. Integrality determines the parity of that valuation, and $3X_t(y) + 1 = 2^{2t+e(y)}y$ proves the claim. There is no odd predecessor when $3 \mid y$. With $R_0 = 0$, one also has $X_t(y) = 2^{e(y)}yR_t + (2^{e(y)}y - 1)/3$. For $q \geq 1$, the multiplier is invertible modulo 3^q , so the root-period proof in Section 2, which also applies to nonnegative indices, shows that every 3^q consecutive t 's give every residue once. Filtering out nonunits retains exactly $2 \cdot 3^{q-1}$ roots. $\square$

Lemma 1.5 (removing a possible cycle predecessor). For any function F and fixed y , at most one z with $F(z) = y$ is periodic. The nonperiodic points in that fiber satisfy the analogue of (W.antichain). A periodic predecessor puts y on its cycle, where there is exactly one cyclic predecessor of y . If $F^j(z) = w$, $j > 0$, for two predecessors, then $F^j(y) = y$ and $w = F^{j-1}(y)$; thus w is periodic. This also rules out recurrence of one nonperiodic predecessor. $\square$

Corollary 1.6 (the complete finite-depth hierarchy). Define

$$\mathcal{C}_0 = \{1\}, \quad \mathcal{C}_{d+1} = \{X_t(y) : y \in \mathcal{C}_d, 3 \nmid y, t \geq 0, X_t(y) > 1\}. \quad (\text{W.depthclasses})$$

Then $\mathcal{C}_d$ is exactly the set of positive odd integers with first accelerated hitting time d to one. Induct on d , using (W.constructor) and uniqueness of the next odd value. Excluding one prevents extending its cycle, and $3 \mid y$ supplies no predecessors. For $d \geq 1$, different members of $\mathcal{C}_d$ cannot lie on one another's forward orbit: such a visit would strictly decrease the remaining first-hitting depth. Thus every finite subset is an antichain. The hierarchy exhausts finite convergent odd trajectories, without asserting that every odd integer belongs to a finite level.

2. Reference transfer and capacity of physical histories

The word formulas now connect actual trajectories to the auxiliary residue distribution of [1, 2]. The affine transfer follows [3, §2], and the selected-family comparison follows [3, §3]. We give their proofs with the word orientation and normalization used here, then bound the total trajectory weight in each residue class. These capacity bounds will compare weighted sums with the reference distribution. The quantitative mixing estimate is stated here and proved in the later mixing section.

The reference law and its normalization

Write $G_q = \mathbb{Z}/3^q\mathbb{Z}$, including the one-point group G_0 , and $\langle f \rangle_q = 3^{-q} \sum_{y \in G_q} f(y)$. A dyadic rational has an unambiguous image in every G_q : a fraction with denominator a power of two is evaluated using the inverse of that denominator modulo 3^q . At positive levels, the units are the residues not divisible by three. We call q the *ternary residue level* of the modulus 3^q . For $n \geq m$, let $\pi_{n,m} : G_n \rightarrow G_m$ be reduction modulo 3^m . On an auxiliary probability space let K_i be independent with $\Pr(K_i = a) = 2^{-a}$, $a \geq 1$, and put $S_i = K_1 + \dots + K_i$. Define

$$C_n = \sum_{i=1}^n 3^{i-1} 2^{-S_i} \pmod{3^n}, \quad \mu_n(y) = \Pr(C_n = y), \quad f_n = 3^n \mu_n, \quad \rho_n = \frac{2}{3} f_n. \quad (\text{T.law})$$

Thus $f_0 = 1$ and $\rho_0 = 2/3$. This normalization is applied once. The function f_n is the probability density relative to the uniform measure on G_n , and ρ_n is its scaled version. A *marker* means the numerical factor

$\rho_n(x \bmod 3^n)$. To mark a trajectory weight, multiply it by this factor at the current end of the inverse construction, which is the source of the corresponding forward trajectory. Its *marked mass* is the sum of those products, whereas *unmarked mass* is the sum of the weights themselves. Markers are not indicator functions. In particular $\langle f_n \rangle_n = 1$ and $\langle \rho_n \rangle_n = 2/3$. For $n \geq 1$ their support consists of units. Reversing the iid word shows that C_n has the same law as n iterations of $y \mapsto (3y + 1)/2^K$ starting at zero. We use the inverse- prefix order because its affine offset is exactly the one in (W.offset).

For $n \geq m$, the first m terms of C_n are C_m and every later term is zero modulo 3^m . Consequently μ_n projects to μ_m . Equivalently, the average of f_n over each fiber of $G_n \rightarrow G_m$ is f_m . This is the projectivity identity for the reference law.

The only analytic premise in this section is a fixed nonnegative envelope ϵ_m satisfying, for every $n \geq m \geq 1$,

$$\langle |f_n - f_m \circ \pi_{n,m}| \rangle_n \leq \epsilon_m. \quad (\text{T.mix})$$

Section 6 specifies the explicit coefficient and mixing envelope used in the density formula. The left side is the unhalved L1 distance after uniform lifting, averaged over the full residue group. For ρ , the distance is multiplied by $2/3$.

One word as an affine injection

Fix an inverse word w of length d and total valuation A . Let $\omega_w = 3^d 2^{-A}$ and C_w be its offset from (W.offset). If $q \geq d$, the map

$$J_{w,q} : G_{q-d} \longrightarrow G_q, \quad J_{w,q}(z) = C_w + 3^d 2^{-A} z \quad (\text{T.injection})$$

is well defined and injective. Its image is the single fiber $y = C_w \bmod 3^d$. Indeed, changing z by 3^{q-d} changes the expression by a multiple of 3^q , and 2^{-A} is a unit. This also proves bijectivity onto that fiber. For a real function g on G_{q-d} , define

$$(\mathcal{T}_{w,q}g)(y) = \begin{cases} \omega_w g(J_{w,q}^{-1}(y)), & y \in \text{im } J_{w,q}, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{T.operator})$$

Lemma 2.1 (exact transfer norms). This operator is linear and positive, and

$$\langle \mathcal{T}_{w,q}g \rangle_q = 2^{-A} \langle g \rangle_{q-d}, \quad \langle |\mathcal{T}_{w,q}g| \rangle_q = 2^{-A} \langle |g| \rangle_{q-d}. \quad (\text{T.norm})$$

Proof. Sum over the affine image using its unique preimage. The scaling is $3^{-q} \omega_w 3^{q-d} = 2^{-A}$. Positivity makes the same calculation valid for absolute values. The empty word gives the identity, including $q = 0$. $\square$

At an odd integer root M , membership in this affine image is equivalent to integrality of the last inverse endpoint. It also implies integrality of every earlier endpoint. More explicitly, the inverse recurrence has denominator 3^i at step i , and the full inverse numerator reduced modulo 3^i is its step- i numerator multiplied by the unit 2^{A-S_i} . Each integral endpoint is odd by induction, since the numerator $2^{k_i} x_{i-1} - 1$ is odd.

Positivity is a separate condition. If $M > C_w$, then $M > C_v$ for each prefix v , because the offset partial sums increase. The identities $x_i = (M - C_{w|i})/\omega_{w|i}$ prove every endpoint positive. Under this condition the value of (T.operator) at M is exactly the physical inverse word weight times its endpoint marker. The condition $M > C_w$ therefore identifies the residue transfer with positive integer trajectories.

Prefix-free stopping and its exact deficit

Let V be a finite prefix-free set of inverse words, each of length at most h . Set $p(V) = \sum_{w \in V} 2^{-A(w)}$ and $q = h + k$, where $k \geq 1$. A length- q iid word can have at most one prefix in V ; therefore $p(V) \leq 1$. Put

$$A_{V,q} = \sum_{w \in V} \mathcal{T}_{w,q} \rho_{q-|w|}, \quad L_{V,q,k} = \sum_{w \in V} \mathcal{T}_{w,q} (\rho_k \circ \pi_{q-|w|,k}). \quad (\text{T.selected})$$

Lemma 2.2 (selected reference density). Pointwise $0 \leq A_{V,q} \leq \rho_q$, and

$$\langle A_{V,q} \rangle_q = \langle L_{V,q,k} \rangle_q = \frac{2}{3}p(V), \quad \langle \rho_q - A_{V,q} \rangle_q = \frac{2}{3}(1 - p(V)). \quad (\text{T.deficit})$$

Proof. Condition the length- q iid word on having prefix w . Its remaining $q - |w|$ letters are independent geometric variables, and its offset is $C_w + \omega_w C_{\text{tail}}$. Its contribution to the full density is therefore $\mathcal{T}_{w,q} \rho_{q-|w|}$. The prefix events are disjoint. Their union gives the selected nonnegative subdensity, proving the pointwise inequality. Its mean and that of L follow from (T.norm) and $\langle \rho_r \rangle = 2/3$. Subtract the means for the last identity. This also covers V empty or containing the empty word. $\square$

Theorem 2.3 (finite stopping transfer). For $1 \leq \ell \leq q$,

$$\langle |L_{V,q,k} - \rho_\ell \circ \pi_{q,\ell}| \rangle_q \leq \frac{2}{3}\{p(V)\epsilon_k + \epsilon_\ell + 1 - p(V)\} \leq \frac{2}{3}\{\epsilon_k + \epsilon_\ell + 1 - p(V)\}. \quad (\text{T.stopped})$$

Proof. Replace each coarse ρ_k by $\rho_{q-|w|}$. The triangle inequality, (T.norm), and (T.mix) bound the replacement cost by $(2/3)\epsilon_k \sum_w 2^{-A(w)}$. The selected density differs from ρ_q by the exact nonnegative deficit in (T.deficit). Finally replace ρ_q by the lifted ρ_ℓ at cost $(2/3)\epsilon_\ell$. The family retains its original reference weights. $\square$

The same statements apply to any fixed subfamily of V . If the family is chosen differently for different physical parents, this theorem alone does not control that choice: partition the parents by their chosen family and pay the comparison for each nonempty group, or prove a different coupling.

Capacity of actual integer histories

Let $\mathcal{H}$ be a finite family of physically realized inverse histories from a fixed positive odd root M . We call a pair consisting of the word length and total valuation a *tag*. At most $T \geq 0$ distinct total tags (d, A) occur. Assume that each history has weight $\omega = 3^d 2^{-A} \leq \sigma$, where $\sigma \geq 0$, offset in $[c_-, c_+]$, where $c_- \leq c_+$, and that within each tag an integer source supports at most one history. Write $C = c_+ - c_-$ and

$$H_q(a) = \sum_{h \in \mathcal{H}: x_h \equiv a \pmod{3^q}} \omega_h.$$

Theorem 2.4 (offset-window capacity). For every $q \geq 0$ and every residue a ,

$$3^q H_q(a) \leq T(C + 3^q \sigma), \quad \left| \sum_a H_q(a) \psi(a) \right| \leq T(C + 3^q \sigma) \langle |\psi| \rangle_q. \quad (\text{T.capacity})$$

Proof. Fix a tag. Its slope ω is constant, and $M = \omega x_h + C_h$ confines all sources to an interval of length C/ω . Sources in a residue class modulo 3^q have spacing at least 3^q . With no duplicate source within the tag, there are at most $1 + C/(\omega 3^q)$ of them. Multiplication by ω bounds their mass by $\sigma + C/3^q$. Sum at most T tags. Pairing the pointwise bound with $|\psi|$ proves the second inequality. $\square$

Any restriction of $\mathcal{H}$ obeys the same bound, without renormalization. If total depth and valuation are both at most S , one may take $T = (S + 1)^2$. For B concatenated blocks of slopes at most β , one may take $\sigma = \beta^B$. The boundary term $3^q \sigma$ is retained at every level.

Pairing (T.stopped) with such a parent histogram yields a physical marked variation bound of

$$\frac{2}{3} T(C + 3^q \sigma) (\epsilon_k + \epsilon_\ell + 1 - p(V)). \quad (\text{T.pairing})$$

To identify this pairing with an orbit extension, check both integrality and positivity as above for every selected child. All comparisons remain exact when the formal residue class contains just one occupied physical source.

A finite root period

Lemma 2.5 (root residues). For $q \geq 0$, every interval of 3^q consecutive positive indices s contains exactly one $R_s = (4^s - 1)/3$ in each residue modulo 3^q . For $q \geq 1$, R_s is a unit exactly when 3 does not divide s .

Proof. For $t \geq 1$, $\nu_3(4^t - 1) = 1 + \nu_3(t)$. If $3 \nmid t$, then $(4^t - 1)/3 = 1 + 4 + \dots + 4^{t-1}$ is t modulo three. Next, if $a \equiv 1 \pmod{3}$, then $a^2 + a + 1$ has valuation one: writing $a = 1 + 3z$ gives $3(1 + 3z + 3z^2)$. Factoring $a^3 - 1$ shows that multiplying t by three raises the valuation by one. This proves the formula for all t . Now $R_{s+t} - R_s = 4^s(4^t - 1)/3$ has valuation $\nu_3(t)$. Two indices have the same root residue exactly when they agree modulo 3^q . Finally $R_s \equiv s \pmod{3}$. $\square$

A finite root family can therefore realize a full ternary average exactly. To use this average uniformly at later stages, we must also control the total subsequent change of the marks at each fixed root.

3. Stopped reference blocks and a global continuation schedule

The preceding identities apply to any suitable finite word family. We now construct families by stopping when the affine slope first becomes small enough. Grouping several stopped words controls both size growth and the ordinary step count. Crucially, these families depend on the prescribed block index, not on the integer to which they are attached. Their probability space is the reference law of (T.law).

A constant-slope stopping rule

Let $\lambda = \log_2 3$ and $\beta = 1/8$. For an inverse word w define its depth d , total valuation A , displacement $Y = A - \lambda d$, and ordinary step count $R = A + d$. Its slope is 2^{-Y} , so the displacement is the negative base-two logarithm of that slope. Let τ be the first positive depth with $\sum_{i \leq \tau} (K_i - \lambda) \geq 3$. Write Y, R also for the stopped totals. For integer $E \geq 1$, let $V(E)$ consist of first-crossing words with $A < E$, and set $p(E) = \sum_{w \in V(E)} 2^{-A(w)}$, $\varepsilon(E) = 1 - p(E)$.

The set $V(E)$ is finite because its positive parts sum to at most $E - 1$, and prefix-free because a first-crossing word cannot properly extend another. It may be empty for small E . Every word in it has slope at most β .

Lemma 3.1 (stopped exponential moments). The stopping time is almost surely finite, and $\mathbb{E}e^{sR} < \infty$ for some $s > 0$. The vector (Y, R) has an exponential moment in a neighborhood of zero.

Proof. Put $X = K - \lambda$. Its mean is $2 - \lambda > 0$. Its MGF is finite around zero, so for some $t > 0$, $q = \mathbb{E}e^{-tX} < 1$. If $\tau > n$, the displacement at depth n is below three. Markov's inequality gives

$$\Pr(\tau > n) \leq e^{3t} q^n. \quad (\text{S.tautail})$$

For small $s > 0$, let $g(s) = \mathbb{E}e^{2s(K+1)}$; choose s with $g(s)q < 1$. On $\tau = n$, $R = \sum_{i \leq n} (K_i + 1)$. By Cauchy-Schwarz and $\Pr(\tau = n) \leq \Pr(\tau > n - 1)$,

$$\mathbb{E}(e^{sR}; \tau = n) \leq g(s)^{n/2} e^{3t/2} q^{(n-1)/2}.$$

Summing the geometric majorant proves the claim. Since $|Y| \leq (1 + \lambda)R$, a sufficiently small joint exponential moment follows. $\square$

Lemma 3.2 (exact stopped means and concentration). Set $\mu_Y = \mathbb{E}Y$ and $\mu_R = \mathbb{E}R$. Then

$$\mu_Y = (2 - \lambda)\mathbb{E}\tau \geq 3, \quad \mu_R = 3\mathbb{E}\tau, \quad \frac{\mu_R}{(\log 2)\mu_Y} = \frac{3}{\log(4/3)}. \quad (\text{S.means})$$

For each fixed $\delta > 0$ there are finite positive C_δ, c_δ such that independent stopped copies obey

$$\Pr\left(\left|\sum_{i=1}^l Y_i - l\mu_Y\right| > \delta l \text{ or } \left|\sum_{i=1}^l R_i - l\mu_R\right| > \delta l\right) \leq C_\delta e^{-c_\delta l}. \quad (\text{S.corridor})$$

Proof. The event $\tau \geq i$ depends only on $K_1, \dots, K_{i-1}$. Independence and absolute integrability allow summing $\mathbb{E}[K_i 1_{\tau \geq i}] = 2\Pr(\tau \geq i)$. Thus $\mathbb{E}\sum_{i \leq \tau} K_i = 2\mathbb{E}\tau$, which gives the two means without an unbounded optional-stopping assertion. The stopped displacement is at least three pointwise. The ratio follows from $(2 - \lambda)\log 2 =$

$\log(4/3)$. For a centered coordinate $Z = Y - \mu_Y$ or $R - \mu_R$, its log-MGF is differentiable at zero with derivative zero. Choose $t > 0$ small enough that $\log \mathbb{E} e^{tZ} < t\delta/2$, and likewise for $-Z$. Markov's inequality bounds each of the four tails by $e^{-t\delta l/2}$; take the minimum exponent and sum the four bounds. $\square$

These constants are effective: (S.tautail) and the displayed geometric majorant give explicit tails for stopped-word expectations and MGFs. At fixed depth, a large-valuation tail is bounded by the geometric MGF; the depth tail is bounded by the majorant above. Finite rational searches for strict inequalities therefore terminate. Their numerical optimization is not used in the density score.

An elementary precision tail with a fully fixed range

Lemma 3.3 (two-binomial-tail certificate). For every integer $E \geq 256$,

$$\varepsilon(E) \leq 2e^{-E/75} \leq 2^{-E/100}. \quad (\text{S.timeout})$$

Proof. Put $N = \lfloor 2E/5 \rfloor$ and $L = \lceil \lambda N + 3 \rceil$. If $L \leq \sum_{i \leq N} K_i < E$, then the displacement has crossed three by depth N and the first-crossing total is below E . Consequently

$$\varepsilon(E) \leq \Pr\{\text{Bin}(L-1, 1/2) \geq N\} + \Pr\{\text{Bin}(E-1, 1/2) \leq N-1\}. \quad (\text{S.binomial})$$

This is the exact waiting-time correspondence: the sum of N positive geometric variables is the trial of the N -th success. Both $E-1$ and $N-1$ in the second event are essential.

For a sum of $n \geq 1$ fair bits, either centered tail at distance $a \geq 0$ is at most $e^{-2a^2/n}$. Indeed a centered bit has MGF $\cosh(t/2) \leq e^{t^2/8}$, since the derivative of $\log \cosh u$ is $\tanh u \leq u$ for $u \geq 0$; symmetry handles $-t$. Independence and optimization at $t = 4a/n$ prove the bound.

Use $\lambda < 8/5$, proved by $3^5 < 2^8$. With $n = L-1$,

$$n < \frac{8}{5}N + 3 \leq \frac{16}{25}E + 3 \leq \frac{2}{3}E, \quad N - \frac{n}{2} > \frac{1}{5}N - \frac{3}{2} \geq \frac{2}{25}E - \frac{17}{10} \geq \frac{1}{15}E.$$

The last inequalities hold for $E \geq 256$. The first probability in (S.binomial) is therefore at most $e^{-E/75}$. In the second event the downward distance is $(E-1)/2 - (N-1) \geq E/10 + 1/2$, so its probability is at most $e^{-E/50}$. Their sum is at most $2e^{-E/75}$. Finally $\log 2 < 7/10$: the first four terms of $e^{7/10}$ sum to more than two. Thus $E(1/75 - (\log 2)/100) \geq 256(19/3000) > \log 2$, proving the second bound. $\square$

This conservative exponent fixes an effective continuation schedule. It does not replace the sharper reference survival certificate used for the initial word family. In the next section we show that changing to this continuation after a sufficiently long initial history loses arbitrarily little marked mass, with that initial choice fixed before the cutoff limit.

Global indices, finite parsing, and bounded offsets

Fix once and for all

$$E_i = \lceil 1000 \log_2(i+2) \rceil \quad (i \geq 1). \quad (\text{S.precision})$$

Then $E_i \geq 256$ and (S.timeout) gives $\varepsilon(E_i) \leq (i+2)^{-10}$. For any word of depth d , (W.offset) and $K_i \geq 1$ give

$$0 \leq C_w \leq \sum_{i=1}^d 3^{i-1} 2^{-i} = (3/2)^d - 1. \quad (\text{S.offset})$$

For concatenated words $w_i \in V(E_i)$, successive slopes are at most β . The composition identity (W.compose) bounds the complete offset by the fixed finite number

$$C_\infty = \sum_{i \geq 1} \beta^{i-1} ((3/2)^{E_i-1} - 1) < \infty. \quad (\text{S.offsetsum})$$

The bracket grows at most polynomially in i , while β^{i-1} is geometric. A ratio bound on a polynomial-geometric tail gives effective truncation. The factors β^{i-1} provide the accumulated contraction in this sum.

Concatenations of a prescribed finite number of prefix-free blocks have a unique parsing. Parse the first word at its first crossing, remove it, and repeat with the next prescribed precision. If one concatenation properly extended another with the same number of component blocks, this parsing would force every component to agree and leave no extension. Thus the concatenation family is prefix-free, as is every restriction of it.

Rational corridors and finite selection

Fix a rational outer tolerance $\delta > 0$. Effective approximation of (S.means) supplies rational centers satisfying

$$|\tilde{\mu}_Y - \mu_Y| < \delta/4, \quad |\tilde{\mu}_R - \mu_R| < \delta/4. \quad (\text{S.centers})$$

For a concatenation of l stopped blocks retain the rational inner bounds

$$|Y - \tilde{\mu}_Y l| \leq (\delta/2)l, \quad |R - \tilde{\mu}_R l| \leq (\delta/2)l. \quad (\text{S.innercorridor})$$

The true $\delta/4$ -corridor is contained in this inner corridor, which is contained in the true δ -corridor. Both inclusions follow from the triangle inequality. Thus selection preserves the outer geometric bounds while its rejected mass is controlled by (S.corridor) at tolerance $\delta/4$.

This is a decidable finite selection. The R -comparison is rational. For a rational endpoint a/b , $b > 0$, comparison of $Y = A - d \log_2 3$ with a/b reduces, when $bA - a \geq 0$, to comparing the integers 2^{bA-a} and 3^{bd} . If $bA - a < 0$, its sign decides the comparison; depth zero is rational. Both inequalities defining the absolute-value bound are handled this way.

For the continuation construction, use the effective constants $C_{\delta/4}, c_{\delta/4}$ from (S.corridor). Every transition and macroblock below uses these concentration constants, while its geometric conclusions use the outer tolerance δ .

Parent-independent macroblocks at the actual current count

Fix $\delta = 1/20000$ for the later clock, choose the centers above, and take a fixed $L > 0$ with $c_{\delta/4}L > 10$. For cumulative original-block count $B \geq 0$, define

$$\ell(B) = \lceil L \log(B+2) \rceil, \quad B_{j+1} = B_j + \ell(B_j). \quad (\text{S.schedule})$$

A *macroblock* groups $\ell(B)$ successive stopped blocks. The next family $\mathcal{C}_B$ consists of concatenations $w_i \in V(E_i)$, $B < i \leq B + \ell(B)$, satisfying the following decidable displacement and step-count corridors:

$$|\sum Y(w_i) - \tilde{\mu}_Y \ell(B)| \leq (\delta/2)\ell(B), \quad |\sum R(w_i) - \tilde{\mu}_R \ell(B)| \leq (\delta/2)\ell(B). \quad (\text{S.macro})$$

It is finite, prefix-free, and independent of the physical parent. Its reference mass p_B obeys

$$1 - p_B \leq \ell(B)(B+3)^{-10} + C_{\delta/4} e^{-c_{\delta/4} \ell(B)} = O((B+2)^{-9}), \quad h_B \leq \sum_{i=B+1}^{B+\ell(B)} (E_i - 1) = O(\log^2(B+2)). \quad (\text{S.macrodeficit})$$

Proof. Independently sample complete stopped blocks. They have total mass one by (S.tautail); restarting after a finite stopping word gives the same law because each concatenation has probability equal to the product of its word probabilities. Reject if any block exceeds its prescribed precision or either inner corridor fails. By (S.centers)–(S.innercorridor), that failure implies failure of the true quarter-tolerance corridor. The union bound, (S.timeout), and (S.corridor) give the estimate. The depth bound follows from $d \leq A < E_i$ for every component. Since $\ell(B) = O(\log(B+2))$ and $c_{\delta/4}L > 10$, the final polynomial bound has constants independent of the starting count B_0 . $\square$

In particular $p_B > 0$ eventually. At a fixed starting B_0 , every retained continuation satisfies the summed corridors, with only its fixed incoming prefix contributing additive constants. The map from displacement and step count to valuation and depth is invertible:

$$d = (R - Y)/(1 + \lambda), \quad A = (\lambda R + Y)/(1 + \lambda). \quad (\text{S.tags})$$

Windows of width $O(B)$ in both Y, R therefore permit $O((B+1)^2)$ total integer tags, with constants depending only on the fixed corridors. Applying (T.capacity) retains this tag cost, the offset envelope, and the boundary term $3^q \beta^B$. At $k = \lceil B/1000 \rceil$ and $q = k_{\text{next}} + h_B$, q/B tends to $1/1000 < \log_3 8$, so that boundary term decays exponentially. For the full history from block one, if $\epsilon_k \leq C_6/k^6$, (T.pairing) has summable increments: the unrestricted $O(B^2 \log^2 B)$ tag bound times $O(B^{-6} + B^{-9})$ is summable even over every integer B . The tighter corridor tag count removes the harmless $\log^2 B$ factor when its literal hypotheses are available.

For the literal application, set $k_{\text{old}} = \lceil B/1000 \rceil$, $k_{\text{new}} = \lceil (B + \ell(B))/1000 \rceil$, and $q = k_{\text{new}} + h_B$. In (T.stopped), use $k = k_{\text{new}}$ and its $\ell = k_{\text{old}}$. The hypotheses $1 \leq k_{\text{old}} \leq q$ hold for $B \geq 1$. Any initial $B = 0$ stage is separate finite bookkeeping; positive-level mixing is used only for the eventual tail above.

When the continuation starts after a large count B_0 , the schedules (S.precision) and (S.schedule) use that actual count. Restarting a fixed small macro index would leave an initial rejection probability without a bound. The next section also retains the number of incoming tags in the error bound and proves that its product with the continuation tail tends to zero as the initial history is extended.

4. A finite graft with uniformly vanishing weighted loss

We need to extend a family with positive marked mass while retaining control of the trajectories' sizes and step counts. An initial family of words is called a *seed family*. A *graft* keeps a finite initial segment of seed blocks, appends a finite transition, and then continues with the stopped macroblocks of the preceding section. Its *splice index* is the number of seed blocks retained. The estimate below bounds the marked mass lost in this change of continuation, uniformly in the root.

Quantitative reference mixing and lower bounds on the seed marks are explicit premises here. They are established in the mixing, seed-mark and persistent-root sections. All estimates below retain the actual inverse-history weights $\omega = 3^d 2^{-A}$.

The seed words and their geometric interface

Fix an integer $b_0 \geq 32^5$, and define

$$b_{t+1} = b_t + \lfloor b_t/100 \rfloor, \quad m_b = \lceil b^{3/5} \rceil, \quad \gamma = \frac{201}{200}, \quad g = \frac{2013}{2000}. \quad (\text{G.seed})$$

Here t counts seed blocks. At stage t , take exactly the words

$$V_t = \{w : |w| = b_t, \quad |A(w) - 2b_t| \leq m_{b_t}\}. \quad (\text{G.fixedseed})$$

The parameter b_t also specifies the parent-size threshold 16^{b_t} . Every family is finite, since its positive letters have bounded total valuation, and prefix-free, since all its words have the same length. The seed fixes the depth and restricts the total valuation. The later stopped blocks retain their own first-crossing rule from [Section 3](#); their survival and graft estimates are established separately below.

For a positive odd root $M \geq 16^{b_0}$, let $\mathcal{P}_t(M)$ contain the physically realized concatenations of the first t such words. The empty concatenation at $t = 0$ is included. Histories retain their complete words and cuts; identical copies are not counted as different histories. Write x_h, ω_h, C_h for source, slope and offset. All integral positive odd sources are admitted, including nonunits; markers at positive level vanish on the latter.

Lemma 4.1 (seed geometry and actual capacity). Set

$$C_{\text{tag}} = 20(4b_0 + 1), \quad C_{\text{in}} = C_{\text{tag}}(2^{b_0+1} + 16^{b_0}), \quad P_t = C_{\text{in}}(t+1)^3 g^t. \quad (\text{G.capacityconstant})$$

Every history in $\mathcal{P}_t(M)$ satisfies

$$x_h \geq 16^{b_t}, \quad \omega_h \leq 16^{b_0 - b_t}, \quad 0 \leq C_h \leq 2^{b_0+1}, \quad b_0 \gamma^t \leq b_t \leq b_0(101/100)^t. \quad (\text{G.prefix})$$

There are at most $C_{\text{tag}}(t+1)^3 g^t \leq P_t$ total depth/valuation tags. For $H_{t,q}^{\text{in}}(a) = \sum_{h: x_h \equiv a \pmod{3^q}} \omega_h$,

$$3^q H_{t,q}^{\text{in}}(a) \leq P_t \quad (0 \leq q \leq 2b_t). \quad (\text{G.incapacity})$$

These bounds are independent of the root and persist under restriction of the history family without renormalization.

Proof. Write $\lambda = \log_2 3 < 8/5$, using $3^5 < 2^8$. For $b \geq 32^5$, $m_b/b \leq 2b^{-2/5} \leq 1/512$, so every selected word satisfies

$$Y = A - \lambda b \geq (2 - \lambda)b - m_b > b/3, \quad C_w \leq (3/2)^b - 1 < 2^b. \quad (\text{G.seedgeometry})$$

The offset estimate is (S.offset). At a parent $R \geq 16^b$, it gives $R - C_w > R/2$, and hence

$$x = \frac{R - C_w}{\omega_w} > 2^{b/3-1} R, \quad \omega_w < 2^{-b/3}.$$

Since $b/3 - 1 \geq 4\lfloor b/100 \rfloor$ on the stated range, $x \geq 16^{\lfloor b/100 \rfloor} R$ and $\omega_w \leq 16^{-\lfloor b/100 \rfloor}$. This also gives positivity for every compatible word and its prefixes. Multiplying proves the endpoint and slope bounds in (G.prefix).

The composition identity (W.compose) bounds the total offset by

$$\sum_{j < t} 16^{b_0 - b_j} 2^{b_j} \leq 2^{b_0} \sum_{j \geq 0} 2^{-3j} < 2^{b_0+1},$$

since $b_j - b_0 \geq j$. Also $\lfloor b/100 \rfloor \geq b/200$ for $b \geq 200$; the recurrence bounds in (G.prefix) follow.

For the tag count put $B_* = \sum_{j < t} b_j$ and $m_* = \sum_{j < t} m_{b_j}$. Total depth is exactly B_* , and total valuation lies in $[2B_* - m_*, 2B_* + m_*]$. There are therefore at most $2m_* + 1$ total depth/valuation tags. The exact comparison $(101/100)^3 < g^5$, together with $m_b \leq 2b^{3/5}$, gives $m_* \leq 2tb_0^{3/5} g^t$, whence

$$2m_* + 1 \leq (4b_0 + 1)(t + 1)g^t \leq C_{\text{tag}}(t + 1)^3 g^t. \quad (\text{G.seedtags})$$

We retain the latter conservative constant in the continuation estimates.

A source and the total depth determine the canonical forward word; the prescribed block lengths determine all seed cuts. There is at most one history per source in each tag. Apply (T.capacity) with offset width 2^{b_0+1} and slope bound $16^{b_0-b_t}$. If $q \leq 2b_t$, its boundary term is at most $16^{b_0}(9/16)^{b_t} \leq 16^{b_0}$, giving (G.incapacity). Restriction can only remove nonnegative weights. $\square$

Inputs and parameter order for the graft

Assume (T.mix) with $\epsilon_k \leq C_6/k^6$ for $k \geq 1$, where the nonnegative C_6 is fixed. The density recipe uses the particular source coefficient to be specified in its mixing section; the present implication holds for any coefficient satisfying this premise. For $t \geq 0$ put $a_t = \lfloor b_t/4 \rfloor \geq 1$, and define the incoming mark

$$Z_t(M) = \sum_{h \in \mathcal{P}_t(M)} \omega_h \rho_{a_t}(x_h \bmod 3^{a_t}). \quad (\text{G.inmark})$$

The geometric lemma above does not assert that this mark is positive. The later seed-score argument is to supply a fixed finite nonempty pool $\mathcal{R}$ of distinct positive odd roots, all at least 16^{b_0} , and numbers $z_M \geq 0$ such that, for some fixed N ,

$$Z_t(M) \geq z_M \quad (t \geq N, M \in \mathcal{R}), \quad 0 < W \leq \sum_{M \in \mathcal{R}} \frac{z_M}{M}. \quad (\text{G.marks})$$

The graft theorem preserves the pool, its persistent lower marks z_M , and its score W . When counting sources we additionally require (W.antichain) and finite paths from the roots to the chosen target.

Fix the stopping section's $\delta = 1/20000$, $\beta = 1/8$, means μ_Y, μ_R , rational centers (S.centers), and concentration constants $C_{\delta/4}, c_{\delta/4} > 0$ at tolerance $\delta/4$. Choose one integer $L \geq 1$ with $c_{\delta/4}L > 10$. The global precisions are $E_i = \lceil 1000 \log_2(i+2) \rceil$, and macro lengths are $\ell(B) = \lceil L \log(B+2) \rceil$. All these choices precede the splice index t , and t will precede any physical counting cutoff.

The transition and its error bound

At a proposed splice t , write

$$E = \lfloor \sqrt{b_t} \rfloor, \quad B_0 = \lfloor E/4 \rfloor, \quad k_0 = \lceil B_0/1000 \rceil, \quad Q = 2b_t. \quad (\text{G.transition})$$

Take B_0 words from $V(E)$, and retain the concatenation only if its total displacement and step count satisfy

$$|Y - \tilde{\mu}_Y B_0| \leq (\delta/2)B_0, \quad |R - \tilde{\mu}_R B_0| \leq (\delta/2)B_0.$$

These decidable inner bounds imply the outer δ -corridors about the true means, by (S.centers)–(S.innercorridor). This fixed transition family is finite, prefix-free and independent of the physical parent. It has maximum depth at most $B_0(E-1) \leq b_t/4$. For $E \geq 256$, its reference mass p_t obeys

$$1 - p_t \leq B_0 2^{-E/100} + C_{\delta/4} e^{-c_{\delta/4} B_0}. \quad (\text{G.transitionfailure})$$

Indeed sample B_0 complete stopped blocks first, of total reference mass one. The union of their precision failures has probability at most $B_0 \varepsilon(E)$; (S.timeout) and the quarter-tolerance concentration bound give the two terms. Retained words keep their original reference weights.

Apply this same family to each incoming endpoint, keeping its original weight. Mark the resulting physical endpoint at level k_0 , and call the total $Z_{t,0}^{\text{hyb}}(M)$. Physical realization for every compatible branch will be established below before using the comparison. For sufficiently large t , $k_0 \geq 1$, $a_t \leq Q$, and the maximum transition depth plus k_0 is at most Q . Use (T.stopped) with the possibly enlarged depth bound $h = Q - k_0$, rather than requiring Q to be the minimal possible level. Then (G.incapacity) gives

$$|Z_{t,0}^{\text{hyb}}(M) - Z_t(M)| \leq D_t^{\text{tr}} := \frac{2}{3} P_t \left\{ \frac{C_6}{k_0^6} + \frac{C_6}{a_t^6} + B_0 2^{-E/100} + C_{\delta/4} e^{-c_{\delta/4} B_0} \right\}. \quad (\text{G.transitiondebit})$$

The bound includes both mixing terms and the rejected reference mass. In particular, the incoming estimate is used only at its allowed level $Q = 2b_t$.

This error bound tends to zero uniformly in the root. To see the dependence explicitly, for $E \geq 256$ one has

$$E \geq \sqrt{b_t}/2, \quad \sqrt{b_t}/16 \leq B_0 \leq \sqrt{b_t}/4, \quad k_0 \geq \sqrt{b_t}/16000, \quad a_t \geq b_t/8.$$

For fixed $c > 0$, $e^{-cu} \leq r! c^{-r} u^{-r}$, by the r -th term of the exponential series. Use $r = 7$ on the precision term, including its factor B_0 , and $r = 6$ on the corridor term. It follows that $D_t^{\text{tr}} \leq C_{\text{tr}} P_t b_t^{-3}$ for a fixed effective C_{tr} . Since $g < \gamma^3$, $(t+1)^3 (g/\gamma^3)^t \rightarrow 0$, as required.

Composed offsets and realization of every retained branch

It is useful to replace the stopped section's infinite offset envelope by a finite rational upper bound. Let

$$\overline{C}_\infty = \sum_{i=1}^{1997} 8^{1-i} (i+2)^{1000} + \frac{4}{3} 8^{-1997} 2000^{1000}. \quad (\text{G.offsetconstant})$$

Then $C_\infty \leq \overline{C}_\infty$. In fact $(3/2)^{E_{i-1}} - 1 < (i+2)^{1000}$. For $D_i = 8^{1-i} (i+2)^{1000}$ and $i \geq 1998$, the binomial theorem gives

$$\frac{D_{i+1}}{D_i} = \frac{1}{8} \left(1 + \frac{1}{i+2} \right)^{1000} \leq \frac{1}{8(1 - 1000/(i+2))} \leq \frac{1}{4}.$$

Summing this geometric tail proves the assertion.

The transition counts as the first B_0 appended stopped blocks. Its exceptional common precision E is not substituted into the later logarithmic schedule. At every later macro start use its actual cumulative appended count:

$$B_{j+1} = B_j + \ell(B_j), \quad \mathcal{C}_{B_j} \text{ uses the indices } B_j < i \leq B_{j+1}. \quad (\text{G.continuation})$$

Every complete appended word, through any such stage, has slope at most β^B and offset at most

$$C_{\text{app}} \leq \frac{(3/2)^{E-1}}{1-\beta} + \bar{C}_\infty.$$

For the transition this is a geometric sum over its B_0 contractions. For the later blocks the preceding contractions are β^{i-1} , including the transition, so their contribution is a tail of the global envelope. The continuation therefore uses its actual cumulative block index.

Compose with an incoming seed history:

$$M = \omega_{\text{pre}} \omega_{\text{app}} x + C_{\text{pre}} + \omega_{\text{pre}} C_{\text{app}}.$$

By (G.prefix), $E \leq b_t$, and $3/2 < 2$,

$$\omega_{\text{pre}} C_{\text{app}} \leq 16^{b_0} \left\{ \frac{8}{7} 2^{-3b_t} + \bar{C}_\infty 2^{-4b_t} \right\} \rightarrow 0. \quad (\text{G.offsetabsorption})$$

Choose a fixed splice threshold so that the displayed bound is at most one thereafter, and put $C_* := 2^{b_0+1} + 1$. Every total offset is then at most $C_* < 16^{b_0} \leq M$, uniformly in root, stage and level. The same bound holds for prefixes of a selected appended word, whose offset partial sums only increase. By (T.injection), compatibility gives integrality of every endpoint; the strict offset margin and (W.inverse) give positivity and oddness. Thus each formal compatible branch counted by the transfer is exactly a physical branch with the stated valuations. These are precisely the branches used in the transition and continuation comparisons.

Uniform continuation capacity and its complete tail

The transition corridor and all subsequent macro corridors add. At a cumulative count $B \geq B_0$, every retained appended word satisfies

$$|Y_{\text{app}} - \mu_Y B| \leq \delta B, \quad |R_{\text{app}} - \mu_R B| \leq \delta B. \quad (\text{G.totalcorridor})$$

Under (S.tags), the possible depths occupy an interval of length $4\delta B/(1+\lambda)$, and the valuations an interval of length $2\delta B$. Since $\lambda > 1$ and $2\delta < 1$, there are at most $(B+1)^2$ appended integer tags. Together with the incoming tag count this permits at most $P_t(B+1)^2$ combined total tags.

Within a combined total tag, a source determines its complete forward word. Its reversal determines the first t seed cuts, then the B_0 transition first crossings and all later prescribed macro cuts. Hence there is at most one history per source in a tag, even if multiple pairs of incoming/appended tags have the same sum. Applying (T.capacity) anew, with the composed offset width and $\omega_{\text{pre}} \leq 1$, gives

$$3^q H_{t,B,q}^{\text{hyb}}(a) \leq P_t(B+1)^2 (C_* + 3^q \beta^B) \quad (q \geq 0). \quad (\text{G.hybridcapacity})$$

Unlike (G.incapacity), this fresh spacing estimate is valid at every level. All its constants were fixed before t .

For definiteness, the following explicit threshold suffices to control the boundary term:

$$B_{\text{cap}} = [400000(L+1)]^2, \quad D_{\text{mac}} = L+1 + C_{\delta/4}, \quad K_{\text{var}} = \frac{8}{3}(C_*+1)(2C_6 1000^6 + D_{\text{mac}}). \quad (\text{G.tailconstants})$$

At a macro start B , put $B' = B + \ell(B)$, $k = \lceil B/1000 \rceil$, $k' = \lceil B'/1000 \rceil$, and

$$h_B = \sum_{i=B+1}^{B'} (E_i - 1), \quad q(B) = k' + h_B.$$

For every $B \geq B_{\text{cap}}$, $q(B) \leq B/2$, and therefore $3^{q(B)} \beta^B \leq 1$. Here is a direct bound with no hidden dependence on the splice. Write $a = L+1$. For $B \geq 2$, $\ell(B) \leq a \log(B+2)$. The inequality $\log x \leq 4x^{1/4}$ for $x \geq 1$ shows $\ell(B) \leq 5aB^{1/4} \leq B$ at the stated threshold. For the corresponding indices, $i+2 \leq 2(B+2)$. Since $\log 2 > 1/2$, one obtains

$$h_B \leq 4000a \log^2(B+2), \quad q(B) \leq B/1000 + 4002a \log^2(B+2).$$

Also $\log^2(B+2) \leq 16\sqrt{B+2} \leq 24\sqrt{B}$. Using $\sqrt{B} \geq 400000a$, the last bound is at most

$$\left(\frac{1}{1000} + \frac{96048}{400000}\right)B = \frac{96448}{400000}B < B/2.$$

The logarithmic inequalities follow from the exponential series and $\log 2 = \int_1^2 x^{-1} dx > 1/2$; no numerical fit is used.

The macro failure estimate (S.macrodéficit) satisfies the explicit bound

$$1 - p_B \leq \frac{D_{\text{mac}}}{(B+2)^9}.$$

Indeed $\ell(B) \leq (L+1)(B+2)$, and $e^{-c_{\delta/4}\ell(B)} \leq (B+2)^{-c_{\delta/4}L} \leq (B+2)^{-10}$. Both marker errors are at most $C_6(1000/B)^6$. Pair (T.stopped) at $k', k, q(B)$ with (G.hybridcapacity); all these are positive marker levels once $B \geq B_{\text{cap}}$. The marked increments satisfy

$$|Z_{t,j+1}^{\text{hyb}}(M) - Z_{t,j}^{\text{hyb}}(M)| \leq K_{\text{var}} P_t B_j^{-4}. \quad (\text{G.increment})$$

This uses $(B+1)^2 \leq 4B^2$, the boundary bound, and the two distinct mixing costs. The current marker in $Z_{t,j}^{\text{hyb}}$ is always $\rho_{\lceil B_j/1000 \rceil}$.

The visited B_j 's are increasing integers. For $B_0 \geq 2$,

$$\sum_{n \geq B_0} n^{-4} \leq \int_{B_0-1}^{\infty} x^{-4} dx = \frac{1}{3(B_0-1)^3} \leq 3B_0^{-3}.$$

Consequently the total continuation error bound, including its first macro, is at most $3K_{\text{var}} P_t B_0^{-3}$. The bound is uniform in the splice.

The graft theorem and the fixed-root quantifiers

Theorem 4.2 (finite graft). At all sufficiently late splices t , the construction above is physically realized and, for every $j \geq 0$ and every positive odd root $M \geq 16^{b_0}$,

$$|Z_{t,j}^{\text{hyb}}(M) - Z_t(M)| \leq \Delta_t, \quad \Delta_t = D_t^{\text{tr}} + 3K_{\text{var}} P_t B_0^{-3} \rightarrow 0. \quad (\text{G.graft})$$

Each hybrid marked sequence has a nonnegative limit, with the same bound. If the fixed pool satisfies (G.marks), then for every $\eta > 0$ there is one finite common splice $t \geq N$ such that

$$\sum_{M \in \mathcal{R}} \frac{Z_{t,j}^{\text{hyb}}(M)}{M} \geq W - \eta \quad (j \geq 0). \quad (\text{G.pool})$$

The same holds for its limit. Neither the root pool nor W depends on η , the stage or any subsequent counting cutoff.

Proof. Choose t beyond the offset threshold, with $E \geq 256$ and $B_0 \geq B_{\text{cap}}$. The realization argument validates the finite transition comparison and every continuation comparison. Sum (G.increment) and (G.transitiondebit) to obtain (G.graft). Absolute summability gives the limit and preserves nonnegativity.

The transition error bound has already been shown to tend to zero. For the other term, (G.prefix) and $B_0 \geq \sqrt{b_t}/16$ imply

$$P_t B_0^{-3} \leq 4096 C_{\text{in}} b_0^{-3/2} (t+1)^3 (g/\gamma^{3/2})^t \rightarrow 0,$$

because

$$\gamma^3 - g^2 = \frac{16263}{8000000} > 0.$$

All thresholds and constants are independent of the root and of the chosen later splice; they are effective from the displayed rules and the effective stopped-law constants.

Finally fix the pool and let $S = \sum_{M \in \mathcal{R}} 1/M > 0$. Choose a single $t \geq N$ with $\Delta_t S < \eta$. Use $Z_t(M) \geq z_M$ and sum (G.graft) with the positive coefficients $1/M$. This proves (G.pool) for the pooled marks. If $\eta < W$, at least one retained family is nonempty at every stage. The choice of t is made before a physical cutoff is introduced. $\square$

State supplied to the terminal construction

For the convergence-to-one specialization, suppose the fixed pool is a convergent antichain satisfying (W.antichain), with specified finite root-to-one paths. This includes (R.pool). For the fixed pool and splice, let

$$A_t^{\max} = \sum_{r < t} (2b_r + m_{b_r}), \quad T_t^{\max} = \sum_{r < t} (3b_r + m_{b_r}).$$

The fixed depths and valuation windows give these upper bounds on incoming valuation and ordinary step count. In particular $\omega_{\text{pre}} \geq 2^{-A_t^{\max}}$. Let $M_- = \min \mathcal{R}$, $M_+ = \max \mathcal{R}$, and $T_{\mathcal{R}} = \max_{M \in \mathcal{R}} \tau_1(M)$, finite by the specified paths. Immediate predecessors of one are a special case. For every retained source at a central count $B = B_j$,

$$\begin{aligned} (\mu_Y - \delta)B + \log_2(M_- - C_*) &\leq \log_2 x \leq (\mu_Y + \delta)B + \log_2 M_+ + A_t^{\max}, \\ \tau_1(x) &\leq (\mu_R + \delta)B + T_t^{\max} + T_{\mathcal{R}}. \end{aligned} \quad (\text{G.terminalstate})$$

Indeed $x = (M - C_{\text{total}})2^{Y_{\text{app}}}/\omega_{\text{pre}}$; use $0 \leq C_{\text{total}} \leq C_* < M_-$, the two slope bounds, and (G.totalcorridor). The ordinary step count is the actual sum $A + d$, and adding the root-to-one path gives the second line. These inequalities are uniform over the whole fixed finite pool, even if some root families are empty.

The terminal construction will use these bounds to obtain sources in a common interval at every sufficiently large scale. It must retain the same histories, weights and one fixed splice throughout.

Connection to the terminal construction

The seed used here is the fixed-depth family (G.fixedseed). The estimates above control the transition and continuation errors uniformly in the root. Reference mixing and the persistent seed-mark premise (G.marks) are the two remaining inputs to this construction; the later sections prove them with the choices needed here. The next step is to place the resulting sources in short intervals at every sufficiently large scale.

5. Terminal sources at every sufficiently large scale

The preceding graft preserves a fixed pool's marked mass while controlling the displacement and ordinary step count of its histories. Here we place the outer sources in a short multiplicative interval at every large scale, compress their terminal overshoot, and count each physical source once. This is the interface used by all the later density bounds.

The shifted first-crossing barrier follows [3, §3]; terminal scale selection, even overshoot compression and the reference fan follow [3, §5]. We prove the probability and integer-trajectory estimates needed here, including the ordinary clock for the compressed sources.

Graft and terminal parameter lookup

The numerical construction below uses $\theta = 1/1000$. The parameterized source theorem allows every fixed $0 < \theta \leq 1/1000$; the table records both choices.

Quantities	Role and order of choice	Definition
$\mathcal{R}, W$	Fixed roots and persistent score, chosen before the clock; counting also uses nonrecurrence	(G.marks), (W.antichain)
$y_*, T_{\mathcal{R}}$	Fixed target and finite maximum root-to-target path length; $y_* = 1$ in the convergence construction	(G.terminalstate), (N.targetfamily)
$\theta, \delta = \theta/20$	Terminal ratio and concentration tolerance; positive and fixed before the splice	(N.everyclock), (N.parammargin)
$L, \ell(B), E_i$	Macro-block length $\ell(B) = \lceil L \log(B + 2) \rceil$ and global word precisions; L depends on δ	(S.schedule), (S.precision)

Quantities	Role and order of choice	Definition
t, B_0, B	Seed splice index, initial appended block count, and current appended block count	(G.transition), (G.terminalstate)
μ_Y, μ_R	Mean binary displacement and ordinary step cost per stopped block	(S.means)
$\tilde{\mu}_Y, \tilde{\mu}_R$	Rational centers for the inner corridor; concentration is used at tolerance $\delta/4$	(S.centers), (S.innercorridor)
$b = \lfloor \theta B \rfloor, k = \lceil \theta B \rceil$	Terminal size and reference level, growing with the stage	(N.stage), (N.everyclock)
w_b, l_b, h_b, e_b, r_b	Crossing window, overshoot scale and range of terminal shifts	(N.parameters)
$u, K = e_b$	Parent-dependent barrier shift and retained overshoot cap	(N.words), (N.stage)
I_B, R_X	Common interval of counting cutoffs and actual source-shell radius	(N.common), (N.everyclock)

Here B counts stopped reference blocks; it is not the number of valuation letters. The seed sizes b_j and the terminal size b are distinct. After the clock and tolerance are fixed, the splice and its schedules stay fixed while B, b, k grow with the counting cutoff.

A finite terminal family and its exact reference deficit

Let $\lambda = \log_2 3$ and $J(v) = \lceil \lambda v \rceil$ for nonnegative integers v , with $J(0) = 0$. For an integer $b \geq 200$, put

$$w_b = \lfloor 3b/5 \rfloor, \quad l_b = b - w_b, \quad h_b = b + w_b, \quad e_b = \lfloor b/100 \rfloor, \quad r_b = 2w_b - J(w_b) - 2e_b. \quad (\text{N.parameters})$$

The integer r_b is nonnegative. More precisely,

$$\frac{229}{1000}b - 2 \leq r_b \leq \frac{3}{5}(2 - \lambda)b, \quad e_b \geq b/200. \quad (\text{N.radius})$$

Indeed $\lambda < 317/200$, as $3^{200} < 2^{317}$. Using $w_b \geq 3b/5 - 1$, $J(w_b) \leq \lambda w_b + 1$, and $e_b \leq b/100$ gives the lower bound with additive loss less than two. The upper bound follows from $J(w_b) \geq \lambda w_b$ and $e_b \geq 0$. The last floor inequality holds for every $b \geq 200$.

For an integer $0 \leq u \leq 2r_b$, define the nondecreasing integer barrier

$$H_{b,u}(s) = 2b + J((s - b)_+) - J((b - s)_+) + u - r_b. \quad (\text{N.barrier})$$

For an inverse word $v = (k_1, \dots, k_d)$, write $A_s = \sum_{i \leq s} k_i$. For an integer cap $K \geq 0$, let $\mathcal{V}(b, u, K)$ consist of words satisfying

$$l_b < d \leq h_b, \quad A_s < H_{b,u}(s) \text{ } (l_b \leq s < d), \quad 0 \leq A_d - H_{b,u}(d) \leq K. \quad (\text{N.words})$$

The miss at l_b is required; the first crossing is the first one after that depth, not necessarily the first one from depth zero. This family is finite, since depth and total valuation are bounded, and prefix-free: a proper extension of a selected word already has its eligible crossing. Its unnormalized reference mass is $p(b, u, K) = \sum_{v \in \mathcal{V}(b, u, K)} 2^{-A_d}$.

Lemma 5.1 (two endpoint tails and overshoot). Uniformly over the indicated integers b, u, K ,

$$1 - p(b, u, K) \leq 2e^{-b/64000} + 2^{-(K+1)}. \quad (\text{N.deficit})$$

Proof. In the geometric reference law, put $S_s = A_s - 2s$. The exact endpoint identities are

$$H_{b,u}(l_b) - 2l_b = 2e_b + u, \quad H_{b,u}(h_b) - 2h_b = u - 2r_b - 2e_b. \quad (\text{N.endpoints})$$

If the lower-endpoint miss or subsequent crossing by h_b fails, then

$$S_{l_b} \geq 2e_b \quad \text{or} \quad S_{h_b} < -2e_b.$$

For the second implication, a miss at l_b with no eligible hit leaves $A_{h_b} < H_{b,u}(h_b) \leq 2h_b - 2e_b$. This is an event on the incoming state and its reference continuation.

The exact waiting-time correspondence used in (S.binomial) gives

$$\begin{aligned} \Pr(S_{l_b} \geq 2e_b) &= \Pr\{\text{Bin}(2l_b + 2e_b - 1, 1/2) \leq l_b - 1\}, \\ \Pr(S_{h_b} < -2e_b) &= \Pr\{\text{Bin}(2h_b - 2e_b - 1, 1/2) \geq h_b\}. \end{aligned} \quad (\text{N.coin})$$

The respective distances from the binomial means are both $e_b + 1/2$. Both numbers of trials are positive, and at most $16b/5$: for the first, $2l_b + 2e_b - 1 \leq (41/50)b + 1 \leq b$; for the second use $h_b \leq 8b/5$. The fair-coin Chernoff bound proved in the stopping section and $e_b \geq b/200$ therefore bound each probability by

$$\exp\left(-\frac{2(e_b + 1/2)^2}{16b/5}\right) \leq e^{-b/64000}.$$

For an eligible first-hit prefix ending at depth $d - 1$, let $a = H_{b,u}(d) - A_{d-1}$. The preceding miss and monotonicity of the barrier imply $a \geq 1$. Conditional on the next valuation being at least a , the probability that it exceeds $a + K$ is exactly $2^{-(K+1)}$. Sum over disjoint first-hit events. The overshoot failure has mass $2^{-(K+1)}$ times the uncapped crossing mass, hence at most $2^{-(K+1)}$. This proves (N.deficit). $\square$

Physical realization, compression, and a sharp shell

For a terminal parent $y \geq 16^b$, every residue-compatible word in $\mathcal{V}(b, u, K)$ is physically realized. In fact (S.offset) gives

$$0 \leq C_v \leq (3/2)^d - 1 < 2^b, \quad C_v/y < \eta_b := 8^{-b}. \quad (\text{N.offset})$$

Here $d \leq 8b/5$ and $(3/2)^8 < 2^5$. Compatibility gives integrality of every prefix by (T.injection); the strict offset margin gives positivity, oddness and the exact valuations by (W.inverse). This applies to all caps, not only to the compressed family. If the parent is a multiple of three, no nonempty inverse word is compatible; its positive-level marker is also zero.

The terminal *overshoot* is the excess of the total valuation over its crossing barrier. The next lemma reduces an arbitrary retained overshoot to zero or one by removing an even number from the last valuation. The resulting source relation leads to an affine family, called the *fan*.

For $j \geq 0$ put

$$F_j(x) = 4^j x + \frac{4^j - 1}{3}, \quad H_k(x) = \sum_{j \geq 0} 4^{-j} \rho_k(F_j(x)). \quad (\text{N.fan})$$

The second series is convergent, since ρ_k is a bounded function on a finite residue group. All its arguments are reduced modulo 3^k . Thus $H_k(x)$ combines all the reference factors in the fan, with their relative weights. Keeping the index of a fan member will preserve the identity of the original history during compression.

Lemma 5.2 (even overshoot compression). A physical word with overshoot $O \leq K$ has a unique image consisting of a word with overshoot at most one and the integer $j = \lfloor O/2 \rfloor$. Its parent, earlier inverse endpoints, shift and depth are unchanged. If x and v are the reduced source and complete physical weight, the original source and weight are

$$x_{\text{old}} = F_j(x), \quad v_{\text{old}} = 4^{-j} v. \quad (\text{N.compress})$$

Proof. Subtract $2j$ from the last valuation. If the preceding total is A_{d-1} , the old last valuation is $H_{b,u}(d) - A_{d-1} + O$, whose first term is at least one. The new valuation is positive and its overshoot is zero or one. Earlier misses and the first-crossing depth are unchanged. Removing an even number preserves last-step admissibility modulo three. The reduced word has the same offset bound (N.offset), so every inverse endpoint remains positive and odd.

At the common penultimate endpoint, the identity $3F_j(x) + 1 = 4^j(3x + 1)$ gives the source formula. The total valuation has fallen by $2j$, giving the weight formula. Raising that last valuation by $2j$ recovers the original word. Thus the map is injective while the fan index is retained. $\square$

Define

$$g_b = 2(1 - \eta_b), \quad L_b(y) = \frac{4^b y}{4^{2r_b} 3^b}, \quad L_b^\sharp(y) = g_b L_b(y), \quad \mathcal{R}_b = \frac{16}{1 - \eta_b}. \quad (\text{N.scales})$$

Lemma 5.3 (source bounds after compression). Every compatible terminal source with overshoot at most one satisfies

$$2^u L_b^\sharp(y) \leq x \leq \frac{8}{1 - \eta_b} 2^u L_b^\sharp(y). \quad (\text{N.shell})$$

Proof. For $d = b + v$, the rounding error $J(v) - \lambda v$ belongs to $[0, 1)$; for $d = b - v$, its negative belongs to $(-1, 0]$. Consequently, with overshoot at most one,

$$\frac{1}{2}(4/3)^b 2^{u-r_b} \leq 2^{A_d} 3^{-d} \leq 4(4/3)^b 2^{u-r_b}.$$

Use $x = (y - C_v)2^{A_d} 3^{-d}$, (N.offset), and the definition of $L_b^\sharp$. $\square$

For every $L_b^\sharp(y) < X \leq 2^{2r_b} L_b^\sharp(y)$, choose the least integer $u = u(y, X)$ with $2^u L_b^\sharp(y) \geq X$. Then

$$1 \leq u \leq 2r_b, \quad X \leq 2^u L_b^\sharp(y) < 2X, \quad X \leq x < \mathcal{R}_b X. \quad (\text{N.selector})$$

The strict lower endpoint explains why the chosen shift is at least one; the reference estimates also include the unused boundary shift zero. The ratio $\mathcal{R}_b$ tends to 16. The offset estimate (N.offset) gives the stated finite-radius correction.

One common interval for the fixed pool

Assume (T.mix) with $\epsilon_k \leq C_6/k^6$, and the admissible pool and persistent marks (G.marks), with the convergent antichain specified at (G.terminalstate). Fix a sufficiently late graft as in (G.graft). Its current appended counts are $B = B_j$, with $B_{j+1} = B_j + \ell(B_j)$. Write

$$b = \lfloor B/1000 \rfloor, \quad k = \lceil B/1000 \rceil, \quad K = e_b, \quad \delta = 1/20000. \quad (\text{N.stage})$$

From (G.terminalstate), there are fixed constants c_-, c_+, c_τ , uniform over the whole finite pool, such that every current endpoint y satisfies

$$(\mu_Y - \delta)B + c_- \leq \log_2 y \leq (\mu_Y + \delta)B + c_+, \quad \tau_1(y) \leq (\mu_R + \delta)B + c_\tau. \quad (\text{N.central})$$

These constants include the chosen finite seed/graft and the root-to-one paths. They do not depend on the terminal cutoff or later stage. Since $\mu_Y \geq 3$, all these endpoints satisfy $y \geq 16^b$ eventually. We always work beyond that threshold and $b \geq 200$.

Put

$$a(B) = \log_2 g_b + (2 - \lambda)b - r_b - 2, \\ A(B) = (\mu_Y + \delta)B + c_+ + a(B), \quad D(B) = (\mu_Y - \delta)B + c_- + a(B) + 2r_b, \quad I_B = (2^{A(B)}, 2^{D(B)}]. \quad (\text{N.common})$$

Every $X \in I_B$ is in the selector interval for every retained endpoint: its lower endpoint dominates all $L_b^\sharp(y)$, and its upper endpoint is below all $2^{2r_b} L_b^\sharp(y)$. The interval (N.common) uses the fixed deterministic bounds for the whole pool.

Lemma 5.4 (coverage of a half-line). At all sufficiently late central stages I_B is nonempty, consecutive intervals overlap, and their union contains every sufficiently large real number. A stage chosen for $X \rightarrow \infty$ necessarily has $B \rightarrow \infty$.

Proof. From (N.radius) and $b \geq B/1000 - 1$,

$$D(B) - A(B) = 2r_b - 2\delta B + c_- - c_+ \geq \frac{179}{500000}B - C \quad (\text{N.room})$$

for a fixed C . To check consecutive overlap, use the deterministic endpoints themselves. If $b' \geq b$, the floor and ceiling formulas give $w_{b'} - w_b \leq (3/5)(b' - b) + 1$, $e_{b'} - e_b \leq (b' - b)/100 + 1$, and $J(w_{b'}) - J(w_b) \leq \lambda(w_{b'} - w_b) + 1$. Thus $|r_{b'} - r_b| \leq 3(b' - b) + 7$. Also $|\log_2 g_{b'} - \log_2 g_b| \leq 1$. With $B' = B + \ell(B)$, we have $b' - b \leq \ell(B)/1000 + 1$, so both $|A(B') - A(B)|$ and $|D(B') - D(B)|$ are $O(\log(B + 2))$. All constants here were fixed with the graft. This logarithmic change is smaller than the positive linear width in (N.room), proving both inequalities needed for consecutive intersection.

Finally (N.radius) gives $a(B) \geq \log_2 g_b + (2/5)(2 - \lambda)b - 2$, which tends to infinity. Both interval ends therefore tend to infinity. Their consecutive intersections make the tail union an interval unbounded above, hence it contains a half-line. Any fixed finite collection of the I_B 's is bounded. Thus selecting an interval containing an increasingly large X forces its stage to increase. $\square$

Controlling the parent-dependent terminal choice

For each $X \in I_B$, append $\mathcal{V}(b, u(y, X), K)$ to each current endpoint y , retaining the complete history. Let $Z_{B,X}^{\text{ter}}(M)$ be its physical mass at root M , marked by ρ_k at the original, uncompressed outer source. Write $Z_B^{\text{hyb}}(M)$ for the current graft mark.

The terminal family depends on y , so a single common-family transfer cannot be applied to all parents at once. Partition them by their selected integer shift. There are at most $2r_b + 1 \leq 2b + 1 \leq B$ groups at the stages under consideration. Each restricted histogram inherits (G.hybridcapacity) without renormalization.

Use the ambient level $q = h_b + k$, child marker k , and old marker k in (T.stopped). At $B \geq 200000$,

$$q \leq \frac{13}{5}b + 1 \leq \frac{13}{5000}B + 1 \leq B/2, \quad 3^q \beta^B \leq 1.$$

The all-level graft capacity is therefore applicable, and every compatible terminal branch is physical by (N.offset). Sum the comparisons only after restricting to the actual shift groups. It follows that

$$\sup_{X \in I_B} |Z_{B,X}^{\text{ter}}(M) - Z_B^{\text{hyb}}(M)| \leq d_B := \frac{2}{3}P_t(B+1)^2(C_*+1)(2r_b+1) \left(\frac{2C_6}{k^6} + 2e^{-b/64000} + 2^{-(e_b+1)} \right) \rightarrow 0. \quad (\text{N.payment})$$

Here P_t, C_* are the fixed graft constants. The limit is uniform in M and X : $k \geq B/1000$, $b \geq B/2000$ eventually, and the right side is $O(P_t B^{-3})$ plus polynomial factors times decaying exponentials. Both marker substitutions, the selected-family deficit, the shift count, and the level boundary are included in this bound.

Aggregate source weights and the terminal fan

Replace each terminal family by its physically realized family with overshoot at most one, retaining the same central histories, parent-dependent shift, and fixed stage. Count each complete reduced history just once, not once for each original that compresses to it. For a reduced history z from root M , denote its weight and source by $v(z)$ and x_z .

At a fixed root M in the convergent antichain, a source determines its complete forward orbit and cannot visit M twice. Reversing its uniquely determined word recovers the seed cuts, the fixed transition stops, and the prescribed subsequent stopped-block and macro cuts. The current parent is then determined, its shift is the selector $u(y, X)$, and the terminal word is fixed. Thus a source supports at most one complete reduced history at this stage.

The affine identity gives $x_z v(z) \leq M$. The pool satisfies (W.antichain), so (W.poolcharge) applies after identifying reduced histories. After division by the root and grouping by actual source, this gives

$$a_x := \sum_{M \in \mathcal{R}} \frac{1}{M} \sum_{z: x_z = x} v(z) \leq \frac{1}{x}. \quad (\text{N.charge})$$

By (N.compress), originals inject into pairs (z, j) . First sum only this finite image, then enlarge to all reduced histories and all $j \geq 0$, using nonnegativity. The normalized original terminal mark satisfies

$$\mathcal{A}_{B,X} := \sum_{M \in \mathcal{R}} \frac{Z_{B,X}^{\text{ter}}(M)}{M} \leq \sum_x a_x H_k(x). \quad (\text{N.fancharge})$$

The fan index is summed only in the marked inequality. The reduced source weights in (N.charge) count each history once at the single selected stage.

Fine precision and the reduced-source clock

For any current endpoint y , the lower end of the common interval gives

$$X > L_b^\sharp(y) \geq \frac{g_b 16^b}{4 \cdot 3^b},$$

because $r_b \leq 2b$. Since $k \leq b + 1$,

$$\frac{3^k}{X} \leq \frac{12}{g_b} (9/16)^b \rightarrow 0. \quad (\text{N.precision})$$

This includes the genuine ceiling case $k = b + 1$. The limit is uniform over all cutoffs in the common interval. Also $k \rightarrow \infty$ with the selected stage.

For every terminal word of cap $K = e_b$, monotonicity of J , $u - r_b \leq r_b$, and the exact radius identity imply

$$A_d \leq 2b + J(w_b) + r_b + e_b = 2b + 2w_b - e_b, \quad A_d + d \leq 3b + 3w_b - e_b \leq \frac{24}{5}b.$$

The family with overshoot at most one has the same last bound, since $e_b \geq 2$. Thus the looser retained coefficient 481/100 also bounds its terminal step count.

The reduced source itself is eventually at least its current parent: from (N.shell), $u \geq 0$, and (N.radius),

$$\log_2 \frac{x}{y} \geq \log_2 g_b + (2 - \lambda)b - r_b - 2 \geq \log_2 g_b + \frac{2}{5}(2 - \lambda)b - 2 > 0.$$

Combining this size bound with (N.central) gives, uniformly over reduced sources and cutoffs,

$$\limsup \frac{\tau_1(x)}{\log x} \leq \frac{\mu_R + \delta + 481/100000}{(\log 2)(\mu_Y - \delta)}. \quad (\text{N.clock})$$

Here the limit runs through the central stages, and fixed prefix/root step counts are already included in c_τ . The denominator is a lower bound on the logarithm of the reduced source, not that of its larger original.

By (S.means), the right side is decreasing in $\mu_Y \geq 3$. At three it is

$$\frac{9}{(3 - \delta) \log(4/3)} + \frac{\delta + 481/100000}{(3 - \delta) \log 2} < \frac{1747479290400}{167532267749} < \frac{10431}{1000}. \quad (\text{N.clockmargin})$$

For a short exact verification, the positive atanh series $\log((1+z)/(1-z)) = 2 \sum_{j \geq 0} z^{2j+1}/(2j+1)$ gives $\log(4/3) > 2(1/7 + 1/(3 \cdot 7^3) + 1/(5 \cdot 7^5))$ and $\log 2 > 2(1/3 + 1/(3 \cdot 3^3) + 1/(5 \cdot 3^5)) > 693/1000$. Substituting these lower bounds gives the displayed rational upper bound; the final cross-multiplied gap is

$$10431 \cdot 167532267749 - 1000 \cdot 1747479290400 = 49794489819 > 0.$$

The series follows by integrating the geometric series on $[0, z]$. Because (N.clock) is uniform, increasing the stage threshold for each fixed finite extension ensures that every later reduced source satisfies the literal ordinary step bound $\tau_1(x) \leq (10431/1000) \log x$.

The all-scale physical source-family theorem

Theorem 5.5. Suppose the mixing and finite-pool premises of (G.marks) hold with certified score $W > 0$, and the pool is a convergent antichain as specified at (G.terminalstate). For every $\varepsilon > 0$ and every fixed $\Lambda > 16$, one can choose one finite common graft and a threshold X_0 such that, for every real $X \geq X_0$, there are a finite set of odd sources S_X , nonnegative weights a_x , a level $k(X)$, and a normalized original terminal mark $\mathcal{A}_X$ satisfying

$$\begin{aligned} S_X &\subseteq [X, \Lambda X) \cap \{x : \tau_1(x) \leq (10431/1000) \log x\}, & a_x &\leq 1/x, \\ W - \varepsilon &\leq \mathcal{A}_X \leq \sum_{x \in S_X} a_x H_{k(X)}(x), & & \text{(N.sourcefamily)} \\ k(X) &\longrightarrow \infty, & 3^{k(X)}/X &\longrightarrow 0. \end{aligned}$$

The fixed pool and W do not depend on ε, Λ, X . The graft and its schedules are fixed before X ; its selected terminal stage and shift families may depend on X .

Proof. Put $a = \min\{\varepsilon/2, W/4\} > 0$. Use (G.pool) to fix one common graft with central normalized marks at least $W - a$ at every later stage. The reciprocal-root sum $S = \sum_{M \in \mathcal{R}} 1/M$ is fixed and finite. By (N.payment), $Sd_B < a$ eventually, so $\mathcal{A}_{B,X} \geq W - 2a \geq W - \varepsilon$ uniformly on I_B . In particular positive original terminal mass is retained.

Increase this fixed stage threshold until the common intervals cover their tail, $\mathcal{R}_b < \Lambda$, and the reduced-source clock holds. For every sufficiently large X , select one of these intervals containing it and take its whole pooled reduced family. Apply (N.selector), (N.charge), and (N.fancharge). The stage must tend to infinity with X , so (N.precision) gives both final limits. All limiting assertions are uniform over the cutoffs in each selected interval. The graft is fixed at the chosen positive error tolerance. $\square$

Clocks arbitrarily close to the stopped-mean ratio

Put $c_0 = 3/\log(4/3)$. The fixed numerical construction above has a parameterized version with the same incoming pool and score.

Theorem 5.6 (parameterized source families). Under the mixing and persistent-pool hypotheses of (G.marks), with the same convergent antichain, for every fixed $c > c_0$, every $\varepsilon > 0$, and every fixed $\Lambda > 16$, the conclusion of (N.sourcefamily) holds with $c \log x$ in place of $(10431/1000) \log x$. More precisely, the same fixed score W and root pool admit sources satisfying

$$\begin{aligned} S_X &\subseteq [X, R_X X) \cap \{x : \tau_1(x) \leq c \log x\}, & 0 &\leq a_x \leq 1/x, \\ W - \varepsilon &\leq \mathcal{A}_X \leq \sum_{x \in S_X} a_x H_{k(X)}(x), & & \text{(N.everyclock)} \\ R_X &< \Lambda, & R_X &\longrightarrow 16, & k(X) &\longrightarrow \infty, & 3^{k(X)}/X &\longrightarrow 0. \end{aligned}$$

One finite graft and its schedules are fixed before all sufficiently large real X . They may depend on c, ε, Λ ; the root pool and W do not.

Proof. We check each estimate affected by the new parameter. Fix a rational $0 < \theta \leq 1/1000$, put $\delta = \theta/20$, choose rational centers (S.centers), and choose concentration constants $C_{\delta/4}, c_{\delta/4}$ at tolerance $\delta/4$, with an integer $L \geq 1$ satisfying $c_{\delta/4}L > 10$. Use the decidable inner selectors (S.innercorridor) in the transition and macroblocks. Retain the stopped law, $\beta = 1/8$, global precisions (S.precision), and seed families. Use the same macro schedule with this L, δ , and replace every appended marker level by $\lceil \theta B \rceil$. All these choices precede the splice.

At a splice t , retain $E = \lfloor \sqrt{b_t} \rfloor$, $B_0 = \lfloor E/4 \rfloor$, and $Q = 2b_t$, and take $k_0 = \lceil \theta B_0 \rceil$. The transition rejection bound (G.transitionfailure) holds with the new concentration constants. Eventually

$$B_0 \geq \sqrt{b_t}/16, \quad k_0 \geq \theta \sqrt{b_t}/16 \geq 1, \quad B_0(E - 1) + k_0 \leq Q.$$

The unchanged incoming capacity is therefore available at Q . The two substitutions and rejected mass in (G.transitiondebit) give $D_t^{\text{tr}} \leq C_{\text{tr}}(\theta) P_t b_t^{-3} \rightarrow 0$, using the same exponential-series bounds as before. The offsets and slopes have not changed, so (G.offsetabsorption) still proves physical realization of every compatible branch at a sufficiently late splice.

The new transition and macro corridors still sum to (G.totalcorridor). Since $2\delta < 1$, the appended tag count is at most $(B+1)^2$, and the full capacity remains (G.hybridcapacity), including its incoming multiplier P_t . At a macro start put $B' = B + \ell(B)$, $k = \lceil \theta B \rceil$, $k' = \lceil \theta B' \rceil$, and $q(B) = k' + h_B$. The same depth bound gives

$$q(B) \leq \theta B + 4002(L+1)\log^2(B+2) \leq B/2 \quad (B \geq [400000(L+1)]^2).$$

Indeed $\theta \leq 1/1000$, so the explicit boundary calculation following (G.tailconstants) applies. With

$$K_{\text{var}}(\theta) = \frac{8}{3}(C_* + 1)(2C_6\theta^{-6} + L + 1 + C_{\delta/4}),$$

the increment is at most $K_{\text{var}}(\theta)P_tB^{-4}$. Consequently the combined transition and continuation error is at most

$$\Delta_t(\theta) = D_t^{\text{tr}} + 3K_{\text{var}}(\theta)P_tB_0^{-3} \rightarrow 0. \quad (\text{N.paramgraft})$$

The crucial incoming-factor estimate is unchanged:

$$P_tB_0^{-3} \leq 4096C_{\text{in}}b_0^{-3/2}(t+1)^3(g/\gamma^{3/2})^t \rightarrow 0, \quad \gamma^3 - g^2 = 16263/8000000 > 0.$$

Thus one sufficiently late common splice meets any prescribed positive pool tolerance. The constants may grow as $\theta \downarrow 0$; only a fixed positive θ is being used in this limit.

Now take $b = \lfloor \theta B \rfloor$, $k = \lceil \theta B \rceil$, and $K = e_b$ at the terminal stage. The state (G.terminalstate) holds with the new δ and fixed prefix constants. Since $\mu_Y - \delta > 4\theta$, every central endpoint satisfies $y \geq 16^b$ eventually. Use the same terminal barriers, radii and common intervals (N.common). Their width is

$$D(B) - A(B) \geq ((458/1000)\theta - 2\delta)B - O(1) = \frac{179}{500}\theta B - O(1). \quad (\text{N.paramroom})$$

Floors and ceilings now give $b' - b \leq \theta\ell(B) + 1$; the changes in both interval endpoints are still $O_\theta(\log B)$. The positive linear width implies consecutive overlap, and both endpoints tend to infinity by the same bound on $a(B)$. These intervals therefore cover a half-line at this fixed θ .

The selected-shift partition still has at most $2b + 1 \leq B$ groups eventually. Since $k \leq b + 1$, its level is $q = h_b + k \leq (13/5)b + 1 \leq B/2$ eventually. Formula (N.payment) therefore holds literally with the new b, k . Its right side is $O_\theta(P_tB^{-3})$ plus polynomial factors times exponentials decaying in θB : use $k \geq \theta B$ and $b \geq \theta B/2$ eventually. It tends to zero uniformly over the fixed pool and every cutoff in the selected interval. This retains both marker costs and the cost of the parent-dependent choice.

Compression, its injection with a retained fan index, and (N.charge) are unchanged. The actual radius is still $16/(1 - 8^{-b}) \rightarrow 16$. The ceiling bound $k \leq b + 1$ gives (N.precision), so $3^k/X \rightarrow 0$ and $k \rightarrow \infty$. The reduced source is still at least its parent eventually. Its terminal ordinary step cost is at most $(481/100)b \leq (481/100)\theta B$, hence uniformly

$$\limsup \frac{\tau_1(x)}{\log x} \leq c_\theta := \frac{\mu_R + \theta/20 + (481/100)\theta}{(\log 2)(\mu_Y - \theta/20)}. \quad (\text{N.paramclock})$$

Using (S.means) and $\mu_Y \geq 3$ yields

$$0 < c_\theta - c_0 = \theta \frac{(1 + c_0 \log 2)/20 + 481/100}{(\log 2)(\mu_Y - \theta/20)} < 3\theta. \quad (\text{N.parammargin})$$

For the strict last bound, $\log(4/3) > 2/7 > 3/11$ implies $c_0 < 11$, and $2/3 < \log 2 < 7/10$. The numerator of the last fraction is less than $1049/200$, whereas three times its denominator exceeds $6 - 1/10000 > 1049/200$. Choose the fixed rational θ also to satisfy $3\theta < c - c_0$. Then $c_\theta < c$, so one eventual threshold gives the literal clock $c \log x$ for every later reduced source.

Finally choose one splice meeting the first half of the desired normalized mark tolerance, using (N.paramgraft), and then a stage threshold meeting the terminal half by (N.payment). Enlarge the latter for the common half-line, radius, size, and clock bounds. The proof of (N.sourcefamily) now gives (N.everyclock) for every real cutoff beyond one threshold. The seed, persistent pool and score were fixed before θ , so none depends on c . $\square$

Sources reaching a fixed target

Theorem 5.7 (fixed-target source families). Let $y_\star > 0$ be a fixed integer. Suppose a nonempty finite odd pool satisfies (G.marks) and (W.antichain), and each root has a specified finite ordinary path to $y_\star$. Retain the same mixing hypothesis as in (N.everyclock). For every fixed $c > c_0$, every $\varepsilon > 0$, and every $\Lambda > 16$, every conclusion of that theorem holds with $\mathcal{G}_c(y_\star)$ in place of $\mathcal{G}_c$:

$$\begin{aligned} S_X &\subseteq [X, R_X X) \cap \mathcal{G}_c(y_\star), & 0 \leq a_x \leq 1/x, \\ W - \varepsilon &\leq \mathcal{A}_X \leq \sum_{x \in S_X} a_x H_{k(X)}(x), & (\text{N.targetfamily}) \\ R_X &\rightarrow 16, & k(X) \rightarrow \infty, & 3^{k(X)}/X \rightarrow 0. \end{aligned}$$

The pool, target and score precede c ; the graft and cutoff may depend on c, ε, Λ . No convergence of $y_\star$ is assumed.

Proof. The seed realization, variation and graft proofs are uniform for all odd roots above 16^{b_0} . Let $T_{\mathcal{R}}$ now be the maximum of the specified root-to-target path lengths. Concatenation gives $\tau_{y_\star}(x) \leq A + d + T_{\mathcal{R}}$ for a word of total valuation A and length d . Thus the second line of (G.terminalstate) and (N.central) holds with $\tau_{y_\star}$ and this finite additive constant; all source-size bounds are unchanged.

The terminal selector, shift partition and two marker error bounds depend on those size bounds and incoming capacity. Their proofs are unchanged. Nonrecurrence gives a unique root visit and hence the same complete valuation word and deterministic cuts. After the same compression identification, (W.poolcharge) supplies (N.charge); the finite-image injection still supplies (N.fancharge).

The reduced source is at least its parent eventually, and its terminal ordinary cost is still bounded by $(481/100)\theta B$. Consequently (N.paramclock) holds for $\tau_{y_\star}$, with precisely the same c_θ . Choose $3\theta < c - c_0$ and apply (N.parammargin). The finite target-path cost disappears in that ratio. All interval overlap, radius and precision statements concern the same constructed sources, so the final assembly proves (N.targetfamily). $\square$

Why the present estimates stop short of the endpoint

Every fixed $\theta > 0$ gives $c_\theta - c_0 > 0$ in (N.parammargin). Taking θ smaller controls this upper-bound gap but does not remove it. Setting $\theta = 0$ would instead give $b = k = 0$, outside the terminal construction, and $\delta = 0$, outside its positive-tolerance concentration estimates. It also removes the positive linear width in (N.paramroom), the parameterized form of (N.room), which makes consecutive cutoff intervals overlap.

The graft constants contain θ^{-6} , and the limiting arguments fix θ before the splice and cutoff. A shrinking-parameter construction would require additional uniform control of these errors, coverage and clock slack. The common lower-density bound for $c > c_0$ does not supply that control or pass automatically to $c = c_0$. The nonquantitative diagonal consequence (D.vanishingclock) does give an unspecified vanishing clock loss; it requires no uniform shrinking-parameter estimates.

From source families to density

The result above gives the precise interval coverage, aggregate source weights and ordinary clock needed for counting. Its mixing and persistent-mark premises are inherited through (G.marks). We now establish those premises, then return to this source family to derive the explicit density bound.

6. Explicit reference mixing from primitive Fourier decay

The construction compares reference densities at different ternary levels. We now bound the error in replacing a finer distribution by the uniform lift of a coarser one. The explicit Fourier-decay theorem proved in Appendix E supplies the input. The Fourier decomposition below gives a bound valid at every finer level; we then optimize its auxiliary parameter. This error bound controls both the continuing seed construction and the final conversion from marked mass to source weight.

The primitive theorem and the quantitative target

Use the reference law μ_n of (T.law). For a function $p : G_n \rightarrow \mathbb{C}$, set

$$\hat{p}(\xi) = \sum_{x \in G_n} p(x) e^{2\pi i \xi x / 3^n}.$$

The Fourier transform is unnormalized. A frequency $\xi \in G_n$ is primitive when $3 \nmid \xi$. We first work with any fixed real $C \geq 1$ satisfying the following input:

$$|\widehat{\mu_n}(\xi)| \leq Cn^{-6409} \quad (n \geq 1, \xi \in G_n, 3 \nmid \xi). \quad (\text{M.primitive})$$

The canonical numerical choice of C and the local theorem asserting this input are given in the final subsection. The proofs up to that instantiation are uniform in every such C .

Write $h = \log 2$, $\gamma = \log(4/3)$, $\lambda = \log 3 / \log 2$, and fix

$$D = C20^{6409}, \quad C_6 = 2D + 2, \quad m_0 = 2^{80}, \quad A = 2314/25. \quad (\text{M.constants})$$

For real $80 \leq v \leq 679/5$, put

$$\begin{aligned} a(v) &= 6407 - \frac{1}{2}hv^2, & q_v &= hv, \\ E_v(x) &= 2Dx^{-a(v)} + \frac{32}{q_v^2 \log x} x^{1-q_v} + 2x^{-q_v} + 2x^{-q_v-1} \quad (x \geq m_0). \end{aligned} \quad (\text{M.widthbound})$$

The target is, for every integer $1 \leq m \leq n$,

$$\text{Osc}(m, n) := \sum_{x \in G_n} |\mu_n(x) - 3^{m-n} \mu_m(x \bmod 3^m)| \leq \epsilon_m, \quad (\text{M.target})$$

$$\epsilon_m = \begin{cases} \min\{2, C_6 m^{-A}\}, & m < m_0, \\ \min\{2, C_6 m^{-A}, \min_{80 \leq v \leq 679/5} E_v(m)\}, & m \geq m_0. \end{cases} \quad (\text{M.envelope})$$

Uniform lifting divides each probability at level m equally among its 3^{n-m} refinements at level n . Thus $\text{Osc}(m, n)$ is the sum of absolute differences from that lift. The *mixing envelope* ϵ_m bounds this error for every finer level at once.

All logarithms in this section are natural. The coefficient C is fixed first. The width v is fixed during each proof, but every estimate has the same displayed cutoff throughout its compact range. Consequently one may optimize v at a fixed coarse level m ; it never depends on a physical source or on a subsequent counting cutoff.

A head selected by its first crossing

Fix $n \geq m_0$, write $L = \log n$, and use independent valuations $\Pr(K_i = a) = 2^{-a}$, $a \geq 1$. Put $S_i = \sum_{j \leq i} K_j$, $S_0 = 0$. For a word $w = (k_1, \dots, k_d)$, define its dyadic offset and its integer cleared offset by

$$C(w) = \sum_{i=1}^d 3^{i-1} 2^{-S_i}, \quad B(w) = 2^{S_d} C(w) = \sum_{i=1}^d 3^{i-1} 2^{S_d - S_i}. \quad (\text{M.offset})$$

Here the S_i are the partial sums of that word. The empty offset is zero. For concatenated words w, t , direct splitting gives

$$C(wt) = C(w) + 3^{|w|} 2^{-S_{|w|}} C(t). \quad (\text{M.split})$$

This is the inverse-prefix orientation of (T.law).

A word is v -typical at ambient length n if every nonempty contiguous subword of length r has valuation sum within $v(\sqrt{rL} + L)$ of $2r$. Let $F_{n,v}$ be the failure of this condition for the full length- n word. Set

$$Q = \lambda n - v^2 L. \quad (\text{M.gatelevel})$$

For $0 \leq k < n$ and $0 \leq l < 2n$, the (k, l) head family consists of typical words of length $d = k + 1$, total $S_d = l$, with $S_k \leq Q < S_{k+1}$. Positivity of the valuations makes this crossing unique. The remaining tail keeps its unrestricted law.

We record explicitly why these families cover every globally typical word and why their residue maps are injective at the common cutoff.

Lemma 6.1 (common scalar margins). Throughout $80 \leq v \leq 679/5$ and $n \geq m_0$, one has $L > 4$ and

$$((3/2)v^2 + v)L \leq (3/200)n, \quad ((5/2)v^2 + v)L \leq (3/10)n. \quad (\text{M.logbudgets})$$

Also $1 < \lambda < 8/5$, $693/1000 < h < 347/500$, and $2/7 < \gamma < 1/3$.

Proof. The logarithmic bounds follow, for example, from the positive series

$$\log \frac{1+t}{1-t} = 2 \sum_{j \geq 0} \frac{t^{2j+1}}{2j+1} \quad (0 < t < 1), \quad (\text{M.logseries})$$

obtained by integrating the geometric series. At $t = 1/3$, ten terms give a lower bound H_- , and the omitted tail is at most

$$\frac{2 \cdot 3^{-21}}{21(1 - 1/9)}.$$

Direct rational comparison gives

$$693/1000 < H_- < h < H_+ := H_- + \frac{2 \cdot 3^{-21}}{21(1 - 1/9)} < 693147181/10^9 < 347/500. \quad (\text{M.logbounds})$$

At $t = 1/7$, the first term gives $\gamma > 2/7$; the upper bound follows from $\log(1 + 1/3) < 1/3$. Finally $3 > 2$ and $3^5 < 2^8$ give the bounds on λ .

For $x > 0$, $\log x \leq \sqrt{x}$. Indeed, with $z = \sqrt{x}/2$, $\log z \leq z - 1$ and $\log 2 \leq 1$ imply this inequality. For $n \geq 2^{80}$, therefore $L \leq \sqrt{n}$ and $\sqrt{n} \geq 2^{40}$. Both coefficients in (M.logbudgets) increase with v . The exact comparisons

$$(3/2)136^2 + 136 < (3/200)2^{40}, \quad (5/2)136^2 + 136 < (3/10)2^{40}$$

prove the result. Also $L \geq 80h > 4$. $\square$

Lemma 6.2 (head location). Every globally typical word has a unique head in the families above. Every nonempty head family satisfies

$$20k \leq 17n, \quad Q < l \leq Q + 2vL < 2n. \quad (\text{M.headlocation})$$

Proof. The second budget implies $v^2L \leq n$, so $Q > 0$. The inequality

$$v\sqrt{sL} \leq s/10 + (5/2)v^2L \quad (\text{M.youngindex})$$

follows by completing a square. Applied with $s = n$, typicality and (M.logbudgets) give $S_n \geq (8/5)n > Q$. Thus there is a unique first crossing at a depth $k + 1 \leq n$, and its prefix is typical.

For a nonempty head, typicality of its prefix of length k and $S_k \leq Q$ give, when $k > 0$,

$$(19/10)k \leq \lambda n + ((3/2)v^2 + v)L < (323/200)n.$$

Hence $20k < 17n$. The case $k = 0$ also satisfies the asserted bound. The last valuation is at most $2 + v(\sqrt{L} + L) \leq 2vL$: use $\sqrt{L} \leq L/2$, $L \geq 4$, and $v \geq 17$. Consequently $l \leq Q + 2vL = \lambda n - v(v - 2)L < 2n$. This proves both coverage and the stated finite index ranges. $\square$

Lemma 6.3 (small cleared offset). Every head satisfies

$$0 < B(w) < 3^n. \quad (\text{M.headsmall})$$

Proof. Its prefix sums obey $S_i \geq 2i - v(\sqrt{iL} + L)$. Therefore, writing $j = i - 1$,

$$B(w) \leq 2^{l-2} \sum_{j=0}^{d-1} \exp\{-\gamma j + hv(\sqrt{(j+1)L} + L)\}.$$

Completing a square once more gives

$$hv\sqrt{(j+1)L} \leq (3\gamma/4)(j+1) + \frac{h^2v^2L}{3\gamma}.$$

The remaining geometric sum is at most $(1 - e^{-\gamma/4})^{-1} < 16 < e^{16}$. For the first inequality, $\gamma/4 > 1/14$ and $e^{1/14} \geq 1 + 1/14 > 16/15$ suffice. With $l \leq \lambda n - v^2L + 2vL$, it follows that

$$B(w) < 3^n \exp\{Q_*(v)L + K_*\},$$

$$Q_*(v) = v^2 \left(-h + \frac{h^2}{3\gamma} \right) + 3hv, \quad K_* = -2h + 3\gamma/4 + 16. \quad (\text{M.headmargin})$$

The bounds in (M.logbounds) give

$$-h + \frac{h^2}{3\gamma} < -\frac{196637}{1500000}, \quad Q_*(v) < -\frac{196637}{1500000}v^2 + \frac{1041}{500}v.$$

The displayed quadratic is decreasing for $v \geq 80$ and is less than -672 at 80 . Also $K_* < 17$. Since $L > 4$, the exponent is negative. Positivity of every summand proves the lower bound. $\square$

Lemma 6.4 (head injection and its second moment). At fixed d, l , the map from a head to $C(w) \bmod 3^n$ is injective. Let $p_{k,l}$ be the nonnegative sub-probability mass function on G_n obtained by mapping the (k, l) head family with its original iid weights. Then

$$\sum_{y \in G_n} p_{k,l}(y)^2 \leq 2^{-l}. \quad (\text{M.headL2})$$

Proof. We first show that $B(w)$, together with d, l , determines the word. For $d = 1$, its only valuation is l . For $d \geq 2$,

$$B(w) - 3^{d-1} = 2^{k_d} B(k_1, \dots, k_{d-1}).$$

The shorter cleared offset is odd: its last summand is odd and every earlier summand is even. Thus k_d is exactly the 2-adic valuation of the positive left side. Divide by 2^{k_d} , reduce the total from l to $l - k_d$, and induct.

Equality of the source residues, multiplied by the common unit 2^l , gives equality of the cleared offsets modulo 3^n . By (M.headsmall) they are equal as integers, and the preceding decoding applies. Each accepted head has probability exactly 2^{-l} . Injection makes every occupied atom of $p_{k,l}$ equal to 2^{-l} ; hence its squared sum is $2^{-l} \sum_y p_{k,l}(y) \leq 2^{-l}$. Empty families give zero. $\square$

Tail levels and the squared Fourier estimate

For any real vector p on G_n , define its uniform fiber average and its oscillation by

$$(\mathbf{A}_{m,n}p)(x) = 3^{m-n} \sum_{y \equiv x \pmod{3^m}} p(y), \quad \mathbf{O}_{m,n}(p) = \|p - \mathbf{A}_{m,n}p\|_1. \quad (\text{M.average})$$

Projectivity makes $\mathbf{O}_{m,n}(\mu_n) = \text{Osc}(m, n)$. The averaging map is positive, mass-preserving and L1-contracting. For real vectors p, t and convolution $(p * t)(x) = \sum_y p(y)t(x - y)$, the following finite identity is the precise use of Fourier cancellation.

Lemma 6.5 (finite collision bound). If $1 \leq m \leq n$ and $|\hat{t}(\xi)| \leq \delta$ at every frequency not divisible by 3^{n-m} , with $\delta \geq 0$, then

$$\mathbf{O}_{m,n}(p * t)^2 \leq \delta^2 3^n \sum_y |p(y)|^2. \quad (\text{M.collision})$$

Proof. The geometric character sum on a fiber shows that the Fourier transform of $\mathbf{A}_{m,n}$ keeps exactly the frequencies divisible by 3^{n-m} , and sets all others to zero. Orthogonality of the same characters gives Parseval

in the present normalization: $\sum_{\xi} |\hat{p}(\xi)|^2 = 3^n \sum_y |p(y)|^2$. Also $\widehat{p * t} = \hat{p} \hat{t}$, by changing variables in the finite double sum. If $z = p * t - \mathbf{A}_{m,n}(p * t)$, these identities and Cauchy–Schwarz give

$$\|z\|_1^2 \leq 3^n \sum_y |z(y)|^2 = \sum_{3^{n-m} \nmid \xi} |\hat{p}(\xi) \hat{t}(\xi)|^2 \leq \delta^2 \sum_{\xi} |\hat{p}(\xi)|^2.$$

This proves the claim, including the empty discarded-frequency set when $m = n$. $\square$

Now assume $9n \leq 10m \leq 10n$ and consider a nonempty (k, l) head family. Its length $d = k + 1$ satisfies $d \leq m$: by (M.headlocation), $m - k \geq n/20 \geq 1$. Set $T = n - d$. Let $t_{k,l}$ be the law on G_n of $3^d 2^{-l} C_T$, using the remaining independent T valuations. The selected full-word submass is exactly

$$c_{k,l} = p_{k,l} * t_{k,l}. \quad (\text{M.convolution})$$

Indeed, (M.split) has this form, and the tail is unrestricted and independent of the head; all heads in this slice have the same d, l .

For a discarded frequency, choose its nonzero representative and put $t = \nu_3(\xi) < n - m$. Multiplication by 2^{-l} preserves this valuation. Since $d \leq m$, one has $t < T$. The character applied to $3^d 2^{-l} C_T$ has conductor 3^r , where

$$r = T - t \geq m - k \geq n/20. \quad (\text{M.tailconductor})$$

Projectivity identifies its expectation with a primitive Fourier coefficient of μ_r , at a unit frequency; hence (M.primitive) yields

$$|\hat{t}_{k,l}(\xi)| \leq C r^{-6409} \leq D n^{-6409}. \quad (\text{M.taildecay})$$

The conductor bound gives the factor 20^{6409} .

Proposition 6.6 (general-width high-regime estimate). For $80 \leq v \leq 679/5$, $n \geq m_0$ and $9n \leq 10m \leq 10n$,

$$\text{Osc}(m, n) \leq 2D n^{-a(v)} + 2 \Pr(F_{n,v}). \quad (\text{M.high})$$

Proof. By (M.collision), (M.headL2) and (M.taildecay),

$$\mathbf{O}_{m,n}(c_{k,l})^2 \leq D^2 n^{-12818} 3^n 2^{-l} < D^2 n^{-12818+lv^2},$$

where the strict crossing $l > Q = \lambda n - v^2 L$ gives $3^n 2^{-l} < n^{hv^2}$. Taking the square root preserves half the entropy exponent, before summing any slices.

The first-crossing head events for different (k, l) are disjoint. Their union includes all globally typical words by (M.headlocation). Consequently

$$\mu_n = \sum_{k < n, l < 2n} c_{k,l} + b, \quad b \geq 0, \quad \sum_y b(y) \leq \Pr(F_{n,v}).$$

There are at most $2n^2$ slices. An empty slice contributes zero. Linearity of averaging, the triangle inequality, and $\mathbf{O}_{m,n}(b) \leq 2\|b\|_1$ prove (M.high). The prefix families have not been normalized by their success masses. $\square$

The reference rejection probability

The preceding decomposition has isolated the rejected mass. We now bound that mass directly under the reference law, without charging all interval lengths at their worst common rate.

Lemma 6.7 (geometric log-MGF at its pole). Put $X = K - 2$ and $c = 1/h$. Define $g(s) = \log \mathbb{E} e^{sX}$ for every real $s < h$. For $0 \leq \theta < h$,

$$g(\theta) := \log \mathbb{E} e^{\theta X} = -\theta - \log(2 - e^{\theta}) \leq \frac{\theta^2}{1 - c\theta}. \quad (\text{M.mgf})$$

For every $\theta \geq 0$, $g(-\theta) \leq \theta^2$.

Proof. Summing the geometric series gives the equality. One has $g(0) = g'(0) = 0$ and $g''(\theta) = 2e^{\theta}/(2 - e^{\theta})^2$. For $H(\theta) = \theta^2/(1 - c\theta)$, $H'' = 2/(1 - c\theta)^3$. Their ratio is

$$R(\theta) = \frac{e^{\theta}(1 - c\theta)^3}{(2 - e^{\theta})^2}.$$

It starts at one and is nonincreasing if

$$F(\theta) = 3c(2 - e^\theta) - (1 - c\theta)(2 + e^\theta) \geq 0.$$

On $[0, h] = [0, 1/c]$, $F(0) = 3(c-1) > 0$, $F(h) = 0$, and $F''(\theta) = -e^\theta(c+1-c\theta) < 0$. Concavity proves $F \geq 0$, hence $g'' \leq H''$. Integrate twice from zero. For the other sign, $g''(-\theta) \leq 2$ for $\theta \geq 0$; the same integration gives $g(-\theta) \leq \theta^2$. $\square$

Lemma 6.8 (two-sided sum bound). For $r \geq 1, z > 0$ and independent copies $X_1, \dots, X_r$,

$$\Pr\{|\sum_{i=1}^r X_i| \geq 2\sqrt{rz} + cz\} \leq 2e^{-z}. \quad (\text{M.chernoff})$$

Proof. For either sign apply the exponential Markov inequality at $\theta = s/(1+cs) < h$, $s = \sqrt{z/r}$. The log-MGF bound gives exponent $-\theta(2\sqrt{rz} + cz) + r\theta^2/(1-c\theta) = -z$. Add the two tails. $\square$

Proposition 6.9 (length-sensitive failure). With $q = vh$ and $L = \log n$,

$$\Pr(F_{n,v}) \leq \frac{16}{q^2 L} n^{1-q} + n^{-q} + n^{-q-1}. \quad (\text{M.failure})$$

Proof. For an interval of integer length $1 \leq r \leq L$, use $z = qL + (q/2)\sqrt{rL}$ in (M.chernoff). Writing $s = \sqrt{r/L} \leq 1$, its threshold divided by L is at most

$$2s\sqrt{q + (q/2)s} + c(q + (q/2)s) \leq v + s(2\sqrt{3q/2} + v/2) \leq v(1+s).$$

The last inequality is equivalent to $v \geq 24h$, valid on the stated range. Thus failure of this interval's typicality has probability at most $2n^{-q}e^{-(q/2)\sqrt{rL}}$. There are at most n starts at each length. The decreasing-integrand comparison and substitution $u = \sqrt{xL}$ give

$$\sum_{r \geq 1} e^{-(q/2)\sqrt{rL}} \leq \int_0^\infty e^{-(q/2)\sqrt{xL}} dx = \frac{8}{q^2 L}.$$

This proves the first term in (M.failure).

For $r > L$, use $z = (q+2)L$. Its threshold is at most $v(\sqrt{rL} + L)$, because

$$v \geq 2\sqrt{q+2} + 2c.$$

To check this uniformly, use $h < 7/10$, $c < 3/2$; $(v-3)^2 > 4((7/10)v+2)$ at $v = 80$, with positive derivative thereafter. Each long interval therefore fails with probability at most $2n^{-q-2}$. Bounding their number by the total $n(n+1)/2$ gives the remaining two terms. The strict failure events are contained in the non-strict tail events just estimated. $\square$

Combining (M.high) and (M.failure) gives

$$\text{Osc}(m, n) \leq E_v(n) \quad (n \geq m_0, \quad 9n \leq 10m \leq 10n). \quad (\text{M.localmixing})$$

From nearby levels to the full compact envelope

For $m \leq r \leq n$, projectivity and uniform lifting give

$$\text{Osc}(m, n) \leq \text{Osc}(m, r) + \text{Osc}(r, n). \quad (\text{M.triangle})$$

To see the normalization, insert $3^{r-n}\mu_r$ between the two vectors defining $\text{Osc}(m, n)$. The second difference is the uniform lift of $\mu_r - 3^{m-r}\mu_m$, whose L1 norm is unchanged.

Every power appearing in (M.widthbound) exceeds 8 throughout the width range. Indeed $a(v) \geq a(679/5) > 8$ by (M.logbounds), and $q_v - 1 \geq 80H_- - 1 > 8$. The positive terms are decreasing in x . Since $(11/10)^8 > 2$,

$$2E_v((11/10)x) < E_v(x) \quad (x \geq m_0), \quad (\text{M.contraction})$$

also for the term with $1/\log x$, which is decreasing.

Theorem 6.10 (all high levels). For every fixed permitted width,

$$\text{Osc}(m, n) \leq E_v(m) \quad (n \geq m \geq m_0). \quad (\text{M.allhigh})$$

Proof. Induct on $n - m$, simultaneously for all pairs. The equal case is zero. Put $r = \lfloor 10m/9 \rfloor$. Since $m \geq m_0 > 90$, $r \geq (11/10)m > m$. If $n \leq r$, (M.localmixing) and monotonicity give $E_v(n) \leq E_v(m)$. Otherwise $m < r < n$: apply (M.localmixing) to (m, r) , the induction hypothesis to (r, n) , and then (M.triangle)–(M.contraction). The result is at most $2E_v(r) \leq 2E_v((11/10)m) < E_v(m)$. The induction uses the geometric sequence of intermediate levels. $\square$

Theorem 6.11 (full quantitative mixing). The target (M.target) with the envelope (M.envelope) holds. In particular

$$\text{Osc}(m, n) \leq C_6 m^{-2314/25} \leq C_6 m^{-92} \leq C_6 m^{-6} \quad (1 \leq m \leq n). \quad (\text{M.powers})$$

Proof. Take $v_* = 6749/50$. Use the ten-term rational bounds H_- and H_+ in (M.logbounds); the two margins below exceed $7/1000$ and $1/1000$, respectively. These give

$$6407 - \frac{1}{2}v_*^2 H_+ > 2314/25, \quad v_* H_- - 1 > 2314/25. \quad (\text{M.exponentguards})$$

For $m \geq m_0$, (M.allhigh) therefore implies

$$\text{Osc}(m, n) \leq m^{-A} \left(2D + \frac{32}{q_v^2 \log m} + \frac{2}{m} + \frac{2}{m^2} \right) < (2D + 2)m^{-A} = C_6 m^{-A}.$$

The bracket after $2D$ is less than two already with $q_{v_*} > 80(693/1000) > 55$, $\log m > 4$, and $m \geq 2^{80}$. For $m < m_0$, both compared vectors are probability masses, so $\text{Osc}(m, n) \leq 2$. On the other hand $C \geq 1$ and $A < 100$ give

$$C_6 > 2 \cdot 20^{6409} > 2^{25637} > 2^{8001} > 2(2^{80})^A.$$

Thus $2 < C_6 m^{-A}$ also in this range.

For each fixed $m \geq m_0$, $v \mapsto E_v(m)$ is continuous on the closed interval $[80, 679/5]$, so it attains its minimum. All its values are bounds for the same $\text{Osc}(m, n)$; take their minimum along with the power bound and the universal bound two. This proves (M.envelope). The two weaker powers follow from $m \geq 1$. $\square$

For the paper’s normalized markers the exact identity is

$$\langle |\rho_n - \rho_m \circ \pi_{n,m}| \rangle_n = \frac{2}{3} \text{Osc}(m, n) \leq \frac{2}{3} \epsilon_m. \quad (\text{M.marker})$$

This discharges (T.mix) for the complete retained envelope, and also its sixth-order specialization used in the finite graft and terminal source sections. The order-92 seed variation bound and the later density optimization use the stronger statements with the same C_6 .

The explicit coefficient

The finite recipe in (E.recipe) specifies our coefficient:

$$C_{\text{primitive}} := C_{6409}^{\text{loc}} \geq 1. \quad (\text{M.canonicalC})$$

Theorem 6.12 (primitive input). This coefficient satisfies (M.primitive).

Proof. Apply (E.primitive) at $B = 6409$. Its law is (T.law), its transform is unnormalized, its frequencies are precisely those not divisible by three, and it covers every positive level. $\square$

Together with the preceding uniform-in- C conversion, this fixes $D = C_{\text{primitive}} 20^{6409}$ and $C_6 = 2D + 2$ for the seed, root-score and density recipes. The coefficient and all startup choices precede the physical counting cutoff and the chosen clock.

For comparisons only, (E.mazur-coefficient) preserves the earlier source integer $C_{\text{primitive}}^{\text{Maz}}$. The adapter in [3] is

$$C_{\text{Maz}} = 2C_{\text{primitive}}^{\text{Maz}} 20^{6409} + 2 + 2^{481} > C_6. \quad (\text{M.sourceadapter})$$

The strict comparison follows from (E.coefficientcomparison). Section 11 compares against this fixed coefficient. Appendix A compares fixed-depth and first-crossing seeds at a common primitive coefficient.

Appendix E uses the pairing and separated-triangle method of [1, Section 7] and [3]. The preceding argument converts its coefficient into the envelope (M.envelope), used by the seed and counting constructions.

7. Seed marks and variation bounds uniform over roots

We return to the seed families defined in the continuation construction. Their reference average can be evaluated exactly, but an average at one stage does not guarantee that any fixed root keeps a positive mark later. Using the preceding mixing estimate, we bound the total change at each root over all subsequent stages. The next section combines this variation bound with positive reference survival to choose the required fixed roots.

The finite mark and its exact normalization

Fix an integer $b = b_0 \geq 32^5$. Use the fixed-depth seed families (G.fixedseed), with

$$b_{j+1} = b_j + \lfloor b_j/100 \rfloor, \quad m_j = \lceil b_j^{3/5} \rceil, \quad a_j = \lfloor b_j/4 \rfloor. \quad (\text{B.parameters})$$

Thus $V_j = \{w : |w| = b_j, |A(w) - 2b_j| \leq m_j\}$. Set

$$p_j = \sum_{w \in V_j} 2^{-A(w)}, \quad D_N = \sum_{j < N} b_j, \quad q_N = D_N + a_N. \quad (\text{B.conductor})$$

Every V_j is prefix-free, so $0 \leq p_j \leq 1$. For $N \geq 0$, concatenate one word from each $V_0, \dots, V_{N-1}$ and call the resulting set $\mathcal{V}_N$. At $N = 0$ it contains the empty word. The prescribed lengths determine all cuts. The family $\mathcal{V}_N$ is prefix-free, of exact length D_N , and independent reference letters give

$$\sum_{w \in \mathcal{V}_N} 2^{-A(w)} = \prod_{j < N} p_j. \quad (\text{B.product})$$

The product is one at $N = 0$, and zero if one chosen family is empty.

On G_{q_N} define the backward seed mark by the finite sum

$$g_N(y) = \sum_{w \in \mathcal{V}_N} \mathcal{T}_{w, q_N}(\rho_{a_N} \circ \pi_{q_N - |w|, a_N})(y). \quad (\text{B.mark})$$

Here $\mathcal{T}$ is (T.operator). For a word of length d and valuation A , its nonzero summand is $3^d 2^{-A} \rho_{a_N}(z \bmod 3^{a_N})$, where $y = C_w + 3^d 2^{-A} z$ in G_{q_N} . Here $q_N - d = a_N$ for every complete seed word, so the payload has exactly the required residual precision. No choice of a physical root occurs in this definition, and no family is normalized by its surviving probability.

Lemma 7.1 (finite mean and support). For every $N \geq 0$,

$$g_N \geq 0, \quad g_N(y) = 0 \ (3 \nmid y), \quad \langle g_N \rangle_{q_N} = \frac{2}{3} \prod_{j < N} p_j, \quad \frac{1}{2 \cdot 3^{q_N-1}} \sum_{3 \nmid y} g_N(y) = \prod_{j < N} p_j. \quad (\text{B.mean})$$

Proof. Positivity follows from (T.operator). The full-group mean of each summand is $(2/3)2^{-A(w)}$ by (T.norm), since lifting the marker does not change its mean. Sum and use (B.product). For a nonempty word, $C_w \equiv 2^{-k_1} \pmod{3}$, so its affine image contains only units. For the empty word, $a_N \geq 1$ and the marker itself vanishes on nonunits. There are exactly $2 \cdot 3^{q_N-1}$ units; converting the full-group mean to the unit-relative mean cancels the factor $2/3$. The calculation is an exact average over a complete residue period. $\square$

A pointwise cap for the seed mark

The iid reversal in (T.law) gives the recursion

$$\rho_{r+1}(y) = 3 \sum_{\substack{a \geq 1 \\ 2^a y \equiv 1 \pmod{3}}} 2^{-a} \rho_r((2^a y - 1)/3 \bmod 3^r). \quad (\text{B.recursion})$$

The argument is well defined for $y \bmod 3^{r+1}$. On the coset $y \equiv 1 \pmod{3}$, the allowed letters are even and their total coefficient is one; on $y \equiv 2 \pmod{3}$, they are odd and their total coefficient is two. There are no allowed letters on the third coset. Consequently

$$\|\rho_r\|_\infty \leq \frac{2}{3} 2^r \quad (r \geq 0). \quad (\text{B.marker cap})$$

Indeed, iterate the factor-two bound starting at $\rho_0 = 2/3$.

For completeness, this bound extends to a finite prefix-free stopping tree without assuming equal leaf depths. Give each leaf a nonnegative payload at most B , and suppose the tree has maximum depth D . At a node of depth d , its remaining weighted transfer is at most $2^{D-d}B$. A leaf has value at most $B \leq 2^{D-d}B$. An internal node is a positive sum over its allowed next letters; the total one-letter coefficient at any residue is at most two by (B.recursion). Downward induction therefore proves the bound at that node. A node cannot be both a leaf and an internal node because the family is prefix-free. Missing branches contribute zero. The induction uses the residual level at each node; the padding in (B.mark) ensures that every leaf payload is defined there.

Apply this argument with $D = D_N$ and $B = (2/3)2^{a_N}$. It gives

$$0 \leq g_N(y) \leq B_{q_N} := \frac{2}{3} 2^{a_N}. \quad (\text{B.cap})$$

Thus the cap is proved pointwise. It is not deduced from the mean.

Realization at physical roots

For every positive odd integer $M \geq 16^b$, let $Z_N(M)$ be the physical mark (G.inmark), with the same complete words and cuts. Then

$$Z_N(M) = g_N(M \bmod 3^{q_N}) \quad (N \geq 0). \quad (\text{B.physical})$$

To prove this, fix a word in $\mathcal{V}_N$. Membership of the root residue in its affine image is, by (T.injection), equivalent to integrality of its full inverse endpoint and implies integrality of all earlier endpoints. At the first block the parent is at least 16^b . The local seed estimate in the proof of (G.prefix) gives a strictly positive offset margin for every compatible word and a new endpoint at least 16^{b_1} . Repeat at successive blocks. All endpoints are positive and odd and have exactly the prescribed valuations; no further root-dependent rejection is necessary. Equivalently, the total offset bound $C_w \leq 2^{b+1} < 16^b \leq M$ gives positivity at every prefix.

Conversely every physical history in $\mathcal{P}_N(M)$ arises from one such compatible complete word. Its cuts have the prescribed fixed block lengths. Formula (T.operator) assigns exactly its actual slope weight and endpoint marker. Summing proves (B.physical), including the empty history at $N = 0$. Nonunit endpoints may be retained in the physical family; their positive-level marker is zero. This changes neither side of the identity.

A fixed-time estimate for the valuation window

Work only on the auxiliary iid probability space. Write $S_B = \sum_{i=1}^B (K_i - 2)$. For $B \geq 32^5$, $m = \lceil B^{3/5} \rceil$ satisfies

$$m \geq 32^3, \quad \frac{m}{B} \leq 2B^{-2/5} \leq \frac{1}{512}, \quad \frac{m^2}{B} \geq B^{1/5}. \quad (\text{B.widthguards})$$

By (M.mgf) and $1/\log 2 < 3/2$,

$$\log \mathbb{E} e^{\theta(K-2)} \leq \frac{\theta^2}{1 - (3/2)\theta} \quad (0 \leq \theta < 2/3), \quad \log \mathbb{E} e^{-\theta(K-2)} \leq \theta^2 \quad (\theta \geq 0). \quad (\text{B.mgf})$$

Independence and Markov's inequality, with $\theta = m/(2B + (3/2)m)$ for the positive tail and $\theta = m/(2B)$ for the negative tail, give

$$\Pr(S_B \geq m) \leq e^{-m^2/(4B+3m)}, \quad \Pr(S_B \leq -m) \leq e^{-m^2/(4B)}. \quad (\text{B.fixedtime})$$

The seed failure $|S_B| > m$ is contained in the union of these two events. In particular

$$1 - p_j \leq 2 \exp\left(-\frac{m_j^2}{4b_j + 3m_j}\right) \leq 2 \exp\left(-\frac{b_j^{1/5}}{4 + 3/512}\right). \quad (\text{B.failure})$$

The positive-tail denominator includes the geometric MGF's $3m$ correction. No stopping-window or overshoot error term is needed for this seed.

The increment before simplifying its constants

All reference mixing below uses the fixed coefficient C_6 from (M.constants), with (M.powers) at order 92. A generic order-six coefficient is not sufficient. At stage j , put $B = b_j$ and

$$\ell = a_j = \lfloor B/4 \rfloor, \quad k = a_{j+1} = \lfloor (B + \lfloor B/100 \rfloor)/4 \rfloor, \quad q = B + k. \quad (\text{B.incrementlevels})$$

Both marker levels are at least one and $k \geq \ell \geq B/8$. For the last inequality, $\lfloor B/4 \rfloor \geq B/4 - 1 \geq B/8$ when $B \geq 8$. Moreover

$$\ell \leq q, \quad q \leq \left(1 + \frac{101}{400}\right) B = \frac{501}{400} B < 2B. \quad (\text{B.paidconductor})$$

Thus the comparison level stays within (G.incapacity)'s allowance.

Define $L = L_{V_j, q, k}$ by (T.selected), and let $H_{j, q}^{\text{in}}$ be the actual incoming histogram of (G.incapacity). The same root-independent family V_j is applied to each parent. Its affine compatibility is exactly physical realization: (G.prefix) puts that parent at least at 16^B , so every compatible extension is positive. Uniqueness of the prescribed cut and multiplication of the slope weights then give the exact signed identity

$$Z_{j+1}(M) - Z_j(M) = \sum_{y \in G_q} H_{j, q}^{\text{in}}(y) [L(y) - \rho_\ell(\pi_{q, \ell} y)]. \quad (\text{B.incrementidentity})$$

There is no conditioning or division by retained mass in this formula. Combining (T.stopped), (G.incapacity) and (M.powers) proves, simultaneously for every positive odd $M \geq 16^b$ and every $j \geq 0$,

$$|Z_{j+1}(M) - Z_j(M)| \leq \frac{2}{3} P_j \left(\frac{C_6}{k^{92}} + \frac{C_6}{\ell^{92}} + 1 - p_j \right). \quad (\text{B.rawincrement})$$

In particular this retains the actual incoming tag factor $20(4b+1)(j+1)^3(2013/2000)^j$. Its capacity coefficient is

$$P_j = 20(4b+1)(2^{b+1} + 16^b)(j+1)^3(2013/2000)^j.$$

The 16^b term bounds the finite-modulus boundary error $3^q \omega$; it has not been discarded as a negligible asymptotic term. The two mixing costs and the valuation-window failure are three separate costs in (B.rawincrement).

An every-generation majorant with the retained coefficient

Put

$$\alpha = 100999/100000, \quad g = 2013/2000, \quad \lambda_{\text{var}} = g\alpha^{-92}, \quad S_{92} = 8 \cdot 465! \cdot 576^{465}. \quad (\text{B.rates})$$

For every integer $B \geq 100000$,

$$B + \lfloor B/100 \rfloor \geq (101/100)B - 1 \geq \alpha B.$$

All floors in (B.parameters) meet this condition, so $b_j \geq b\alpha^j$. There is no omitted early-stage exception.

Since $4 + 3/512 < 2333/100$, (B.failure) is at most $2e^{-b_j^{1/5}/(2333/100)}$. For $B \geq 1$, use $2333/100 < 576$, followed by the degree-465 exponential-series term, to get

$$2e^{-B^{1/5}/(2333/100)} \leq 8Be^{-B^{1/5}/576} \leq 8 \cdot 465! 576^{465} B^{1-465/5} = S_{92} B^{-92}. \quad (\text{B.polynomialfailure})$$

Also $465! < 512^{465}$ and $576 < 1024$, whence $S_{92} < 2^{8838}$. The actual coefficient in [the reference-mixing section](#) satisfies $C_6 > 2^{25637}$, so $S_{92} < C_6$. This absorption uses the size of the coefficient; an arbitrary mixing hypothesis without that size condition would not suffice.

Let

$$H_b^* = \frac{2}{3}(4b+1)22(2^{b+1} + 16^b).$$

Replacing the smaller tag coefficient 20 by 22 and using $k, \ell \geq b_j/8$ gives

$$|Z_{j+1}(M) - Z_j(M)| \leq H_b^*(j+1)^3 \frac{2C_6 8^{92} + S_{92}}{b^{92}} \lambda_{\text{var}}^j. \quad (\text{B.onerate})$$

The incoming factor g^j is included in λ_{var}^j . Exact rational comparisons give

$$\frac{2}{5} < \lambda_{\text{var}} < \frac{41}{100} < \frac{1}{2}. \quad (\text{B.ratebounds})$$

Also $4b + 1 \leq 5b$ and $2^{b+1} \leq 16^b$, so $H_b^* \leq 220b16^b$. Finally

$$220(2 \cdot 8^{92} + 1) < 2^{467}. \quad (\text{B.coefficientguard})$$

Indeed $8^{92} = 2^{276}$, $2 \cdot 8^{92} + 1 < 2^{278}$ and $220 < 2^8$, giving even the upper bound $2^{286} < 2^{467}$. We retain the displayed conservative coefficient for the startup recipe; optimizing that variation coefficient is a separate question.

Theorem 7.2 (seed variation uniformly over roots). With exactly

$$F_b = 2^{467}b16^b(C_6 + 1),$$

every positive odd $M \geq 16^b$ satisfies, at every $j \geq 0$,

$$|Z_{j+1}(M) - Z_j(M)| \leq F_b(j+1)^3 \lambda_{\text{var}}^j. \quad (\text{B.variation})$$

Proof. In (B.onerate), use $b^{-92} \leq 1$, $S_{92} < C_6$, $H_b^* \leq 220b16^b$ and (B.coefficientguard), then $C_6 < C_6 + 1$. All bounds are independent of the root and valid starting at $j = 0$. $\square$

Consequently $Z_j(M)$ converges to a nonnegative limit, and for all integers $n \geq N \geq 0$,

$$|Z_n(M) - Z_N(M)| \leq F_b \sum_{j \geq N} (j+1)^3 \lambda_{\text{var}}^j. \quad (\text{B.tail})$$

The series converges by $0 < \lambda_{\text{var}} < 1/2$; sum finite increments first and then bound by its entire nonnegative tail. The same bound holds for the limit. Nonnegativity of that limit does not imply it is positive.

Persistent marks for the root construction

The exact finite mark, cap, mean and physical identity are (B.mark)–(B.physical). The full later-generation loss is (B.tail). To obtain (G.marks), the next step must still certify positive infinite reference survival, choose a finite N for which the entire variation tail fits the retained budget, and select one fixed finite convergent root pool. **The persistent-root section** retains the exact residual (R.residual) and minimum score (R.score); the full variation tail supplies persistence at the fixed roots.

The positive-level markers vanish on nonunit terminal children, which have no further integral inverse child. This permits all positive odd histories to be included in (B.physical). The geometric moment estimate and the mixing bound of Section 6 give the variation bound used below.

Coefficient scope. The same proof of (B.variation) and (B.tail) applies with C_6 replaced throughout by any fixed integer $\widehat{C} > 2^{8838}$ for which (T.mix) holds with

$$\widehat{\epsilon}_m = \min\{2, \widehat{C}m^{-2314/25}\}. \quad (\text{B.coefficientscope})$$

Indeed this supplies both the order-92 and order-six estimates, and $S_{92} < 2^{8838} < \widehat{C}$ supplies the coefficient in (B.onerate). The remaining bounds are the same scalar inequalities. With $\widehat{F}_b = 2^{467}b16^b(\widehat{C} + 1)$, the survival theorem in the next section is unchanged, its variation tail still tends to zero, and its finite startup and root-pool recipes give a positive score. The startup inequalities use $t = \lceil \log_2 \widehat{F}_b \rceil \geq 4b$, which is retained.

The sufficient condition $\widehat{C} > 2^{8838}$ is used through $S_{92} < \widehat{C}$ in (B.onerate). The startup bound $t \geq 4b = 2^{27} > 25637$ holds for every $\widehat{C} \geq 1$. Thus the finite startup argument applies throughout this coefficient range; the explicit construction takes $\widehat{C} = C_6$.

8. Infinite survival and a fixed persistent root score

The preceding section bounded the total future variation of the seed marks. We now show that their reference average stays above that loss after a fixed initial stage. A complete finite period of known convergent roots realizes the reference average. Allocating the remaining mass among these roots gives a positive lower bound for their reciprocal-root weighted sum, called the *root score*. This establishes (G.marks), the premise needed by the finite continuation and terminal-source constructions. The survival estimate below uses a summable fixed-time tail and small rational inequalities.

A short proof of reference survival

Fix $b = 32^5 = 2^{25}$, retaining the fixed-depth stage families, success masses p_j , floors b_j and marks of Section 7. Put

$$\alpha = \frac{100999}{100000}, \quad r = \frac{504}{503}, \quad u = \frac{511}{64}, \quad d_j = 2e^{-ur^j}. \quad (\text{R.deficits})$$

The recurrence gives $b_j \geq b\alpha^j$, and the exact comparisons

$$100999 \cdot 503^5 - 100000 \cdot 504^5 = 32325056940257 > 0, \quad \frac{32}{4 + 3/512} = \frac{16384}{2051} > \frac{511}{64} \quad (\text{R.survivalguards})$$

show that (B.failure) implies $0 \leq 1 - p_j \leq d_j$.

The function $f(t) = e^{-ur^t}$ is decreasing on $[0, \infty)$. Thus $\sum_{j \geq 0} f(j) \leq f(0) + \int_0^\infty f(t) dt$. Substitution $x = ur^t$, followed by $\log r > 1/504$, gives

$$\sum_{j \geq 0} d_j < 2e^{-u} \left[1 + 504 \left(\frac{1}{u} - \frac{1}{u^2} + \frac{2}{u^3} \right) \right]. \quad (\text{R.survivalsum})$$

Here $\log(1+x) > x/(1+x)$ for $x > 0$ follows by integration of $1/(1+t)$. For the integral estimate used in (R.survivalsum), write

$$\int_u^\infty \frac{e^{-x}}{x} dx = e^{-u} \int_0^\infty \frac{e^{-t}}{u+t} dt \leq e^{-u} \left(\frac{1}{u} - \frac{1}{u^2} + \frac{2}{u^3} \right).$$

The pointwise inequality $1/(u+t) \leq 1/u - t/u^2 + t^2/u^3$ holds for $t \geq 0$; multiplication by $u+t$ leaves the nonnegative remainder t^3/u^3 . Integrating $1, t, t^2$ against e^{-t} gives $1, 1, 2$.

The required numerical comparisons are:

$$\frac{1}{u} - \frac{1}{u^2} + \frac{2}{u^3} = \frac{15142976}{133432831} < \frac{227}{2000}, \quad e^u > 2920. \quad (\text{R.exponentialguard})$$

For the second use $u = 8 - 1/64$,

$$e^8 > \sum_{k=0}^{16} \frac{8^k}{k!} = \frac{70233453593}{23648625}, \quad e^{-1/64} > 63/64,$$

and

$$63 \cdot 70233453593 - 2920 \cdot 64 \cdot 23648625 = 5252536359 > 0.$$

Consequently

$$\sum_{j \geq 0} (1 - p_j) < \frac{2}{2920} \left(1 + 504 \frac{227}{2000} \right) = \frac{14551}{365000} < \frac{1}{25}. \quad (\text{R.failurebudget})$$

For a finite product, $\prod_{j < N} p_j \geq 1 - \sum_{j < N} (1 - p_j)$, by induction from $(1-a)(1-b) \geq 1-a-b$ for $0 \leq a, b \leq 1$. The finite products decrease in $[0, 1]$, so their limit exists. We have therefore proved

$$\boxed{\Pi = \frac{24}{25}, \quad \prod_{j < N} p_j > \Pi \ (N \geq 0), \quad \prod_{j \geq 0} p_j \geq \Pi.} \quad (\text{R.certificate})$$

This proves survival from the displayed rational inequalities. The finite reference-moment tables are separate inputs to the quantitative density formula.

Complete variation tail and earliest positive startup

For the canonical score use C_6 from [the reference-mixing section](#). The qualitative alternative (B.coefficientscope) may replace it throughout this section by $\widehat{C}$, retaining both the order-92 estimate and its scalar size condition. A coefficient supplying only order six does not suffice. Define

$$F_b = 2^{467} b 16^b (C_6 + 1), \quad t = \lceil \log_2 F_b \rceil, \quad L = \lceil \log_2 (t + 1) \rceil,$$

$$a = 2013\,100000^{92}, \quad d = 2000\,100999^{92}, \quad \lambda = a/d, \quad V(n) = 2^t T_3(n, \lambda), \quad T_3(n, z) = \sum_{j \geq n} (j+1)^3 z^j. \quad (\text{R.variationdata})$$

Thus t is an integer, $F_b \leq 2^t$, $t \geq 4b = 2^{27}$, and $t < 2^L$. The canonical arithmetic specification of C_6 makes t an effective integer; it is not necessary to print its expansion. By (B.tail), for every admissible root M and every $j \geq n$,

$$|Z_j(M) - Z_n(M)| \leq V(n). \quad (\text{R.latermarks})$$

For $0 < z < 1$, expanding $(n+1+h)^3$ and summing the four absolutely convergent geometric moment series gives

$$T_3(n, z) = z^n \left[\frac{(n+1)^3}{1-z} + \frac{3(n+1)^2 z}{(1-z)^2} + \frac{3(n+1)z(1+z)}{(1-z)^3} + \frac{z(1+4z+z^2)}{(1-z)^4} \right]. \quad (\text{R.exacttail})$$

The geometric moments follow by differentiating $\sum_{h \geq 0} z^h$ on any compact subinterval of $(0, 1)$, or directly by multiplying their convergent sums by $1-z$. Hence this is an exact rational function at $z = \lambda$. Also $V(n) - V(n+1) = 2^t(n+1)^3 \lambda^n > 0$.

The rational inequalities (B.ratebounds) give $2/5 < \lambda < 41/100$. The additional finite integer comparison is

$$2^{131} 2013^{100} 100000^{9200} < 2000^{100} 100999^{9200}, \quad \lambda^{100} < 2^{-131}. \quad (\text{R.hundredrate})$$

If $n \geq 16$ and $0 < z \leq 41/100$, successive terms of $T_3(n, z)$ have ratio at most $(41/100)(18/17)^3 < 1/2$; the last inequality is $478224 < 491300$. Therefore

$$T_3(n, z) < 2(n+1)^3 z^n. \quad (\text{R.tailmajorant})$$

For the fixed-budget comparison and a finite upper bound on the startup, put $N_{\text{fixed}} = \lceil 100(t+3L+214)/131 \rceil$. For all integers $t \geq 25637$,

$$17 \leq N_{\text{fixed}} \leq (100t + 300L + 21530)/131 \leq 4t/5. \quad (\text{R.searchguard})$$

For the upper bound write $k = \lfloor \log_2 t \rfloor \geq 14$, so $L = k + 1$. It suffices that $1500(k+1) + 107650 \leq 24 \cdot 2^k$. At $k = 14$ this is $130150 \leq 393216$; its right increment exceeds its left increment 1500 for every later k . The ceiling bound and the elementary lower bound on N_{fixed} give the other claims.

The hundred-step comparison applies to every integer exponent: $(\lambda^n)^{100} < 2^{-131n}$ for $n \geq 1$. No divisibility condition on n is present. Since $\lambda^{-1} < 5/2$, (R.searchguard) and (R.tailmajorant) imply

$$\begin{aligned} V(N_{\text{fixed}} - 1) &< 2^{t+1} N_{\text{fixed}}^3 \lambda^{N_{\text{fixed}}-1} \\ &< 2(4/5)^3 (5/2) 2^{t+3L-131N_{\text{fixed}}/100} \\ &\leq (64/25) 2^{-214} < 3 \cdot 2^{-214} < \Pi. \end{aligned} \quad (\text{R.searchendpoint})$$

Consequently the exact definitions

$$N = \min\{n \geq 16 : V(n) < \Pi\}, \quad P = \Pi - V(N) > 0 \quad (\text{R.residual})$$

make sense, with $16 \leq N \leq N_{\text{fixed}} - 1$. Strict monotonicity of V permits rational binary search on this finite interval. All values below 16 fail the positivity predicate: the tail then contains $17^3 \lambda^{16} > 17^3 2^{-32}$, and multiplication by 2^t already exceeds one. This defines the startup by a finite integer recipe. The exact value $\Pi = 24/25$, and then the whole residual P , are retained.

For later quantitative comparisons one needs more than $P > 0$. Let $e = d - a > 0$, and for any integer $n \geq 0$ define

$$U_n = d[(n+1)^3 e^3 + 3(n+1)^2 a e^2 + 3(n+1)a(d+a)e + a(d^2 + 4ad + a^2)].$$

Substitution into (R.exacttail) gives $T_3(n, \lambda) = a^n U_n / [d^n (d-a)^4]$. If $P(n) = \Pi - V(n) > 0$, its numerator over the common denominator $25d^n (d-a)^4$ is a positive integer. Because $d < 2^{11+17.92} = 2^{1575}$,

$$P(n) \geq [25d^n (d-a)^4]^{-1} > 2^{-1575n-6305} > 2^{-1575n-7324}. \quad (\text{R.denominator})$$

Here $25 < 2^5$; the final weaker bound is retained for the startup comparison. No reduced-denominator calculation or lower bound on the numerical distance to a threshold is assumed. The integer 2^t adds no denominator.

The fixed finite convergent root pool

Write $q = q_N$ as in (B.conductor), and let

$$H = 2b + 1, \quad \mathcal{S}_q = \{s \in \mathbb{Z} : H \leq s < H + 3^q, 3 \nmid s\}, \quad R_s = (4^s - 1)/3, \quad K_q = 2 \cdot 3^{q-1}, \quad B_q = (2/3)2^q. \quad (\text{R.pool})$$

Here $q \geq \lfloor b/4 \rfloor = 2^{23}$. **The reference-transfer section's** root-period lemma proves that $s \mapsto R_s \bmod 3^q$ is a bijection from $\mathcal{S}_q$ onto the unit residues. Thus this is a finite pool of exactly K_q distinct roots. Each is odd and $R_s \geq R_H = (4 \cdot 16^b - 1)/3 > 16^b > 1$. Moreover $3R_s + 1 = 4^s$, so its next accelerated odd iterate is one, and its ordinary stopping time is $1 + 2s$. The root pool and its maximum clock are finite and fixed before any physical counting cutoff.

Set $g_s = g_N(R_s \bmod 3^q)$. By (B.mean), (B.cap), (B.physical), and the reference survival bound,

$$0 \leq g_s \leq B_q, \quad Z_N(R_s) = g_s, \quad \sum_{s \in \mathcal{S}_q} g_s = K_q \prod_{j < N} p_j \geq K_q \Pi. \quad (\text{R.rootmean})$$

The unit-relative mean on the right is $\prod p_j$, not $(2/3) \prod p_j$. The latter is the full-group mean before restricting to unit support. This step is an exact finite period, not equidistribution of actual sources.

Define the persistent lower marks by

$$z_s = (g_s - V(N))_+.$$

They satisfy $0 \leq z_s \leq B_q$, and since $(x - y)_+ \geq x - y$,

$$\sum_s z_s \geq K_q P, \quad Z_j(R_s) \geq z_s \quad (j \geq N, s \in \mathcal{S}_q). \quad (\text{R.persistence})$$

For the second claim use (R.latermarks) and $Z_j \geq 0$. The bound includes the entire future variation tail. No assumption that every root survives is made.

Exact weighted allocation and discharge of the graft input

For $q \geq 1$ and $0 \leq p \leq 1$, define

$$W(q, p) = \min \left\{ \sum_{s \in \mathcal{S}_q} \frac{v_s}{R_s} : 0 \leq v_s \leq B_q, \quad \sum_s v_s \geq K_q p \right\}, \quad W_* = W(q_N, P). \quad (\text{R.score})$$

The fixed constant W_* names the score for the chosen construction: N is the earliest stage with positive guaranteed residual, and the full residual $P = \Pi - V(N)$ is retained. The minimum accounts for the least favorable allocation of that residual among the specified roots.

The constant H remains fixed. The feasible region is nonempty (take $v_s = p$, since $B_q \geq 4/3$), compact, and finite dimensional. The objective is continuous, so the minimum exists. Its coefficients are strictly

positive; a minimizer has total mass exactly $K_q p$, by decreasing some coordinate if necessary. For $p > 0$, $W(q, p) \geq K_q p / \max_s R_s > 0$.

For the following evaluation only, relabel the roots in decreasing order as $R_1 > \dots > R_{K_q}$. Write $h = \lfloor K_q p / B_q \rfloor$, $r_0 = K_q p - h B_q$. Since $p \leq 1 < B_q$, $h < K_q$. Then

$$W(q, p) = B_q \sum_{i=1}^h R_i^{-1} + r_0 R_{h+1}^{-1}. \quad (\text{R.greedy})$$

Indeed, if a smaller root has positive mass and a larger root has spare capacity, moving the minimum of those two amounts to the larger root strictly decreases the objective. Repeating this exchange proves the displayed formula, including zero mass, exact cap boundaries, and one partial cap. At fixed q , $W(q, p)$ is continuous and strictly increasing in p : between cap boundaries its slope is one of the positive numbers K_q / R_i .

For later comparisons, decreasing the level also has a precise effect on the minimum. At fixed H , $q \geq 2$ and $0 < p \leq 1$,

$$W(q+1, p) < 3 \cdot 4^{-(3^q+1)} W(q, p). \quad (\text{R.conductormonotone})$$

Indeed, allocate mass $3p$ to each unit index in the last interval $[H + 2 \cdot 3^q, H + 3^{q+1})$ of $\mathcal{S}_{q+1}$. It contains exactly K_q such indices; the allocation is feasible, since $B_{q+1} \geq 3$ and $3K_q = K_{q+1}$. Its smallest root exceeds 4^{3^q+1} times the largest root in $\mathcal{S}_q$: the index gap is at least $3^q + 1$, and $R_{s+\Delta} > 4^\Delta R_s$. Compare this feasible upper bound with $W(q, p) \geq K_q p / \max_{\mathcal{S}_q} R_s$. Iterating gives, for $Q > q \geq 2$,

$$W(q, p) > (4/3)^{Q-2} 3^{3^Q-3^q} W(Q, p). \quad (\text{R.conductorgain})$$

This comparison uses the exact cap and minimum-allocation definition, rather than assuming that a smaller root pool automatically helps.

By (R.persistence), the actual lower vector (z_s) is feasible in (R.score). Thus

$$0 < W_* \leq \sum_{s \in \mathcal{S}_q} \frac{z_s}{R_s}, \quad Z_j(R_s) \geq z_s \quad (j \geq N). \quad (\text{R.graftinput})$$

This is precisely (G.marks), with one fixed finite root pool, one fixed startup, and the same explicit score W_* . The score is a lower certificate obtained by minimum allocation; it is not the unknown actual weighted sum of the marks and need not equal that sum.

The fixed-pool graft theorem (G.pool) supplies the input to (N.sourcefamily) with ordinary clock $(10431/1000) \log x$. Its parameterized version (N.paramgraft)–(N.everyclock) supplies, for every fixed $c > c_0$, the same score, aggregate source-weight bound and coverage at every large real scale with clock $c \log x$. The order is: canonical constants, certificate, startup, pool and score; then c , a positive terminal ratio, a positive graft tolerance and one finite splice; finally every sufficiently large cutoff. No root or seed precision is selected afresh as a function of that cutoff. The terminal stage and level may vary along the fixed graft schedule, as in (N.sourcefamily).

General residue-complete blocks and their common score

Let $\mathcal{F}$ be a nonempty finite collection of disjoint root blocks. Fix $q \geq 2$. Each block has one odd root $M_i(r) \geq 16^b$ in every unit residue $r \bmod 3^q$. Put

$$\mathcal{U}_q = (\mathbb{Z}/3^q\mathbb{Z})^\times, \quad A_{\mathcal{F}}(r) = \sum_i \frac{1}{M_i(r)}, \quad \mathcal{V}(q, p) = \{v : 0 \leq v_r \leq B_q, \sum_r v_r \geq K_q p\},$$

$$W_{\mathcal{F}}^{\text{joint}}(q, p) = \min_{v \in \mathcal{V}(q, p)} \sum_r v_r A_{\mathcal{F}}(r), \quad W_i(q, p) = \min_{v \in \mathcal{V}(q, p)} \sum_r \frac{v_r}{M_i(r)}. \quad (\text{R.jointscore})$$

Here $0 < p \leq 1$; at $p = 0$ every score is zero. Each minimum exists, is positive and is rational for rational p .

For the persistent-mark application, specialize to the selected startup N and put

$$\boxed{q = q_N, \quad p = P.} \quad (\text{R.markapplication})$$

Retain the corresponding B_q, K_q . Then $z_r = (g_N(r) - V(N))_+$ is feasible by the mean, capacity and full variation bounds (B.mean), (B.cap) and (R.latermarks), as in (R.persistence). Thus the joint score is an

available persistent score whenever the union has the required counting and target-path properties. The finite allocation statements below remain valid for every $q \geq 2$ and $0 \leq p \leq 1$.

Proposition 8.1 (shared allocation and coverage).

$$W_{\mathcal{F}}^{\text{joint}}(q, p) \geq \sum_i W_i(q, p), \quad (\text{R.jointdominance})$$

with equality exactly when all block objectives have a common minimizing residue allocation. In particular

$$\min_r A_{\mathcal{F}}(r) \geq a_{\mathcal{F}} \implies W_{\mathcal{F}}^{\text{joint}}(q, p) \geq K_q p a_{\mathcal{F}}. \quad (\text{R.coverage})$$

Proof. The feasible set is nonempty and compact. Since every cost is positive, a minimizer uses total mass exactly $K_q p$. The exchange argument of (R.greedy) fills the smallest costs first and proves rationality. All blocks use one common residue vector. For every feasible vector each summand is at least its separate minimum; equality in their sum requires equality in every summand. The coverage bound follows by multiplying its coefficient bound by the total mass. $\square$

Lemma 8.2 (affine block comparison). Suppose a block uses the baseline indices $\mathcal{S}_q$, is residue-complete, and has

$$M_s = \alpha R_s + \frac{\alpha - 1}{3}, \quad 0 < \alpha < 1. \quad W_M(q, p) > \alpha^{-1} W(q, p) \quad (p > 0). \quad (\text{R.blockcomparison})$$

Indeed $M_s < \alpha R_s$ pointwise. Evaluate the strict reciprocal inequality at a minimizer for this block; at least one coordinate has positive mass. The baseline-root objective of that vector is at least $W(q, p)$. This argument uses the common cap and total mass and does not assume invariance of g_N under the residue permutation. A proved upper coefficient $\alpha \leq \beta < 1$ gives the weaker factor β^{-1} . These finite comparisons will be used in Appendix C.

Quantitative score bounds and the fixed-recipe optimum

For $p > 0$, every root in $\mathcal{S}_q$ is less than $4^{H+3^q}/12$. Hence the required total mass gives

$$W(q, p) > 8p 3^q 4^{-(H+3^q)}. \quad (\text{R.scorelower})$$

For $q \geq 3$ and $0 < p \leq 1$, the upper half-period of indices $H + (3^q + 1)/2 \leq s < H + 3^q$ contains at least $3^{q-1} - 1$ indices not divisible by three. This follows either by counting residues modulo three in that interval or by deleting at most one excess multiple of three from its two-thirds proportion. Its total capacity is at least $(3^{q-1} - 1)(2/3)2^q \geq 2 \cdot 3^{q-1} = K_q$: $(1 - 3^{1-q})2^q \geq 3$ already holds at three and increases. Every root there exceeds 2^{3^q} , because $H \geq 1$ and $2s \geq 2H + 3^q + 1$. A feasible allocation supported there proves

$$W(q, p) \leq 3^q 2^{-3^q}. \quad (\text{R.scoreupper})$$

For $p = 0$ this bound follows from $W(q, 0) = 0$. The bound concerns the minimum, not every feasible vector. The right side decreases for integer $q \geq 1$, as its successive ratio is $3 \cdot 2^{-2 \cdot 3^q} < 1$. At the actual $q \geq 2^{23} > 18$, use $3^q < 2^{2q}$ and $3^{18} - 36 = 387420453$ to obtain

$$0 < W_* < 2^{-134217727} < 2^{-4096}. \quad (\text{R.scoredomain})$$

This supplies the small-score domain of the fractional-moment density comparison without evaluating the score or using an empirical estimate.

It is also useful to retain why this particular startup is optimal within the fixed certificate and rounded one-rate majorant. For the level $q(n) = q_n$,

$$q(n) \geq bn, \quad H \leq 8q(n) + 7, \quad q(n+1) \geq q(n) + 1. \quad (\text{R.conductorguards})$$

The first follows by retaining $b_j \geq b$ in its sum, the second from $q(n) \geq \lfloor b/4 \rfloor$, and the third from the new positive summand and nondecreasing terminal quarter. Fix any later $n > N$ and any $0 < p \leq 1$, and

put $q = q_N$. By (R.scoreupper) and its monotonicity, the later score is at most $3^{q+1}2^{-3^{q+1}}$. Combining (R.scorelower), (R.denominator) at N , and (R.conductorguards), gives

$$\begin{aligned} \log_2 \frac{W_*}{W(q(n), p)} &> 3^q - 2H + 3 - \log_2 3 - 1575N - 7324 \\ &> 3^q - 17q - 7337 \geq 12193. \end{aligned} \quad (\text{R.startupoptimum})$$

Here $1575N \leq q$ since $b > 1575$. The last function equals 12193 at $q = 9$ and has positive successive increments $2 \cdot 3^q - 17$; all actual levels exceed nine. Earlier startups leave no positive guaranteed budget under this same majorant. At N , (R.residual) is the largest budget it allows, and fixed-level score monotonicity shows that no smaller budget improves the score. Every later startup loses even if allowed budget one. Thus the selected startup and its full residual uniquely optimize this minimum-score recipe. This does not optimize the underlying variation coefficient, reference family, actual marks, or alternative root pools. The rational residual condition and large-level condition are essential: an arbitrary tiny real budget in an earlier small pool need not beat a later pool.

The coefficient dependence of the startup, residual, and root score is proved in [Appendix B](#).

The score used in counting

The fixed-depth survival proof, complete variation bound and earliest positive allocation provide the score used below. They use the affine transfer, root-period, capacity and finite-graft estimates proved in the preceding sections. Appendix A compares this score with the central first-crossing benchmark.

The reproducibility archive contains the rational survival and rate checks. The all-stage estimates use the primitive coefficient (M.canonicalC); (B.coefficientscope) gives the variant with a general mixing coefficient. We next estimate the reference moments used to convert the score into an ordinary counting bound.

9. Reference moments and their finite certificates

We now bound the concentration of the reference fan by its maximum, variance, and average power of order $3/2$. The counting section will use these three bounds to turn a guaranteed marked sum into a lower bound on total source weight. All moments concern the auxiliary law (T.law). A product inequality propagates finite moment bounds to every level, so finite certificates suffice for the estimates used at arbitrarily large moduli.

Orientation and a conditional affine mixture

Let $T_a(x) = (3x + 1)/2^a$, interpreted in the appropriate ternary ring. The composition $T_{a_n} \cdots T_{a_1}(0)$ is

$$\sum_{i=1}^n 3^{n-i} 2^{-(a_i + \cdots + a_n)}.$$

Reversing a finite iid word turns this into the offset sum in (T.law), without changing its probability $2^{-\sum a_i}$. Thus the forward random affine construction and the inverse-prefix construction have the same law at every level. The physical histories retain the orientation fixed in (W.inverse).

For $p \geq 1$ put

$$M_p(n) = \langle f_n^p \rangle_n, \quad A_n = \max_{G_n} f_n. \quad (\text{F.norms})$$

All averages in this section are over the full group, including zero and the nonunits. We have $M_1(n) = 1$, $M_p(0) = A_0 = 1$.

Theorem 9.1 (moment inheritance). For integers $u, v \geq 0$ and real $p \geq 1$,

$$M_p(u + v) \leq M_p(u)M_p(v), \quad A_{u+v} \leq A_u A_v. \quad (\text{F.inheritance})$$

Proof. Zero-length blocks give equality. Otherwise condition on the last v forward valuations, a word w with probability θ_w . Its affine action is $b_w + 3^v c_w x$, where c_w is a ternary unit. For $y \in G_v$, write the lifts as $y + 3^v z$, $z \in G_u$. For each word with $b_w \equiv y \pmod{3^v}$, the equation for the final residue determines the

preceding residue as an affine permutation $x = \psi_w(z)$ of G_u . Independence of the two reference blocks gives the exact identity

$$f_{u+v}(y + 3^v z) = 3^v \sum_{w: b_w \equiv y \pmod{3^v}} \theta_w f_u(\psi_w(z)). \quad (\text{F.mixture})$$

The nonnegative coefficients sum to $f_v(y)$. If this sum vanishes, the whole fiber vanishes. Otherwise divide by it and apply Jensen's inequality to the probability mixture. Every ψ_w preserves the full average, so averaging first over z and then over y gives the first assertion. Bounding each f_u by A_u instead gives $f_{u+v}(y + 3^v z) \leq A_u f_v(y)$, proving the second. Countably many words cause no problem: the functions on each finite group are bounded, and the coefficients have finite sum; Jensen follows by truncation and one remainder atom, or directly for a countable probability distribution. $\square$

For a probability law on a finite group, its collision probability is the sum of the squared point probabilities: this is the chance that two independent samples agree. In particular the normalized collision energy is $M_2(n) = 3^n \sum_y \mu_n(y)^2$, and is submultiplicative. General p matters: interpolation between the maximum and this second moment does not supply the smaller direct $p = 3/2$ table below.

The general convex-order inequality and the divergence consequences of moment inheritance are collected in [Appendix D](#). The moment inequality above is all that is needed for the finite certificates and their propagation below.

The whole affine fan and the exact coset split

Conditioning on the last forward valuation gives

$$f_{n+1}(y) = 3 \sum_{\substack{a \geq 1 \\ 2^a y \equiv 1 \pmod{3}}} 2^{-a} f_n((2^a y - 1)/3 \bmod 3^n). \quad (\text{F.transfer})$$

The divisibility condition and the modulus 3^{n+1} on y make the argument well defined. Define

$$F_j(x) = 4^j x + (4^j - 1)/3, \quad H_m(x) = \sum_{j \geq 0} 4^{-j} \rho_m(F_j(x)), \quad \phi_m(x) = (3x + 1)/2 \pmod{3^{m+1}}. \quad (\text{F.fan})$$

Every F_j permutes G_m ; the series converges absolutely. The map $\phi_m : G_m \rightarrow G_{m+1}$ is a bijection onto the $2 \pmod{3}$ coset: equality of its values is equivalent to equality of its inputs modulo 3^m . At its image the allowed exponents in (F.transfer) are $a = 2j + 1$, with predecessor $F_j(x)$. Therefore

$$H_m(x) = \frac{4}{9} f_{m+1}(\phi_m(x)). \quad (\text{F.fanidentity})$$

If $y \equiv 1 \pmod{3}$, shifting its even admissible exponents down by one proves $f_n(2y) = 2f_n(y)$ for $n \geq 1$. The nonunit coset is zero. Since multiplication by two bijects the two unit cosets, the p -power sums on them have ratio $1 : 2^p$. Consequently

$$\langle H_m^p \rangle_m = \frac{3(8/9)^p}{1 + 2^p} M_p(m + 1), \quad \langle H_m \rangle_m = \frac{8}{9} =: \mu, \quad \langle H_m^2 \rangle_m = \frac{64}{135} M_2(m + 1). \quad (\text{F.fanmoments})$$

These constants use full-group averages, not unit-relative ones. The maximum likewise lies in the 2-coset, since every value on the 1-coset is doubled on its image. Thus

$$\|H_m\|_\infty = \frac{4}{9} A_{m+1}. \quad (\text{F.fanmax})$$

The word with all $m + 1$ valuations equal to one has probability $2^{-(m+1)}$ and endpoint $-1 \pmod{3^{m+1}}$. Since $\phi_m(-1) = -1$, it also proves the coherent lower bound

$$H_m(-1) \geq \frac{2}{3} (3/2)^m. \quad (\text{F.spike})$$

This all-depth family rules out a polynomial fan maximum. It does not rule out a polynomial collision or fractional moment, because its support occupies a decreasing fraction of the residue group.

The finite inequalities used by the density recipe

The finite certificate consists precisely of

$$A_d \leq a_d \quad (0 \leq d \leq 16), \quad M_2(d) \leq e_d \quad (0 \leq d \leq 16), \quad M_{3/2}(d) \leq b_d \quad (0 \leq d \leq 16), \quad (\text{F.finitetarget})$$

with the following rational values. The first three maximum and collision values are exact; the subsequent comparisons may be proved strictly.

d	a_d	e_d	b_d
0	1	1	1
1	2	5/3	1276143/1000000
2	22/7	15/7	711317/500000
3	24/5	521/200	769803/500000
4	73/10	3069/1000	32779/20000
5	11	707/200	862279/500000
6	17	4001/1000	449901/250000
7	25	2233/500	116627/62500
8	38	1233/250	962777/500000
9	57	5399/1000	1979303/1000000
10	85	2933/500	253519/125000
11	127	3167/500	259093/125000
12	948/5	6803/1000	422721/200000
13	1896/5	7273/1000	134447/62500
14	20856/35	7743/1000	109287/50000
15	3201/5	4107/500	2217669/1000000
16	4802/5	1737/200	561799/250000

The horizons fix the finite work needed by this recipe. Depth sixteen supplies the retained second-moment block rate $e_{16}^{1/16}$; depth twelve supplies the direct fractional rate $b_{12}^{1/6}$ and the rational inequality $b_{12}^2 < (57/50)^{12}$, together with its remainder bounds. The latter rate is smaller, so the comparison compensates for the fractional branch's extra residual-mass power. The quantitative comparison (Q.paid) checks this on its stated small-score range. These are sufficient certificate choices, not minimal depths for a rate advantage: already $b_{11}^{32} < e_{16}^{11}$. The twelve-step fractional formula is retained as a named benchmark. The sixteen-step certificate gives the smaller propagated denominator rate

$$b_{16}^{1/8} = 1.10650942918 \dots < b_{12}^{1/6} = 1.13284538025 \dots \quad (\text{F.sixteenrate})$$

These are rates of certified ceilings, not assertions about the true asymptotic moments. Section 11 compares the complete optimized formulas.

Here is a complete specification of the finite verification and why it is an upper-bound certificate. Let $S = 2^t$, and suppose the nonnegative integer array z bounds Sf_{d-1} entrywise. Positivity of (F.transfer) allows its input to be replaced by z . All divisions below round upward.

In the cyclic algorithm the powers of two enumerate the units modulo 3^d in a cycle of length $\ell = 2 \cdot 3^{d-1}$. Indeed, the valuation formula proved for (T.law)'s root period gives order 3^{d-1} for four and hence order $2 \cdot 3^{d-1}$ for two. Write $y_i = 2^i \bmod 3^d$, and put $g(i) = z_{(y_i-1)/3}$ on the 1-coset and zero otherwise. The infinite sum

$$h(i) = \sum_{a \geq 1} 2^{-a} g(i+a) = \frac{1}{2} (g(i+1) + h(i+1)) \quad (\text{F.cyclic})$$

is periodic. Its first T terms at zero plus $(\max z)2^{-T}$ bound $h(0)$. Starting with that upward-rounded bound, iterate (F.cyclic) backwards at $i = \ell-1, \dots, 1$. Induction gives an upper bound at every index; no approximate fixed-point solve is used. Assign $3h(i)$ to the unit y_i , and zero to the nonunits. The code checks the whole cycle and uniqueness explicitly.

The separate direct algorithm evaluates every unit independently. Let $e = 1$ on the 2-coset and $e = 2$ on the 1-coset, and $x_0 = (2^e y - 1)/3 \bmod 3^{d-1}$, $x_{j+1} = 4x_j + 1$. For J terms put $A = \sum_{j=0}^{J-1} 4^{J-1-j} z_{x_j}$. The exact positive infinite parity sum is bounded by

$$\frac{3A + \max z}{2^{e+2J-2}}. \quad (\text{F.direct})$$

To verify the tail, multiply $3 \cdot 2^{-e} \sum_{j \geq J} 4^{-j} \max z$ by the displayed denominator; the product is exactly $\max z$. This proves both tail and normalization before rounding.

For the depth-sixteen maximum and collision table, use $t = 48$, $T = 96$, $J = 16$, seeding both algorithms with upward rounding of the exact depth-two vector

$$f_2 = (0, 8/7, 16/7, 0, 11/7, 4/7, 0, 2/7, 22/7). \quad (\text{F.exacttwo})$$

Compute $\max z/S$ and $\sum z_y^2/(3^d S^2)$ at every depth through sixteen. Checked unsigned 64-bit arrays store the entries; all transfer, energy and comparison arithmetic uses unbounded integers. Every array assignment rejects overflow. The largest group has $3^{16} = 43,046,721$ entries; no residue class is sampled.

For the fractional rows through twelve, both algorithms also run with $t = 128$, $T = 168$, $J = 32$, starting from $f_0 = 1$. For any even t and upper vector $z \geq 2^t f_d$,

$$M_{3/2}(d) \leq \frac{\sum_y z_y \lceil \sqrt{z_y} \rceil}{3^{d/2} 2^{3t/2}}. \quad (\text{F.sqrtcertificate})$$

The ceiling square root is found by integer square root and a perfect-square test. Thus no floating approximation of a fractional power decides a table inequality. Positive partial sums rounded downward supply independent lower controls. The cyclic $t = 48$, $T = 96$ arrays give the additional fractional rows through sixteen by the same inequality, now with denominator $3^{d/2} 2^{72}$. The final full-group computation gives

$$M_{3/2}(16) \leq \frac{456815260910074290115278246270}{203282392447840896882957090816} < \frac{561799}{250000} = 2.247196. \quad (\text{F.sixteencertificate})$$

The denominator is exactly $3^{16} 2^{72}$. This calculation uses all 43,046,721 residues, checked unsigned 64-bit storage, and unbounded integers for arithmetic and the square-root sum. It requires no second full direct fractional run through sixteen.

As a calibration, an exact rational parity sum through a full period 3^{d-1} computes every entry through depth five; its denominator contains $4^{3^{d-1}} - 1$. It checks (F.exacttwo), full means, coset ratios, projection and the fan identities against both upper algorithms. The implementations and complete certificate records are supplied in the [reproducibility archive](#). The product inequality (F.inheritance) extends the finite table bounds to all levels.

All-level moment bounds

For $n = 16s + r$, $0 \leq r < 16$, and separately $n = 12t + j$, $0 \leq j < 12$, define

$$D_n = e_{16}^s e_r, \quad D_n^{(3/2)} = b_{12}^t b_j. \quad (\text{F.blockceilings})$$

Also define a finite dynamic minimum recursively by

$$\widehat{A}_0 = 1, \quad \widehat{A}_n = \min_{1 \leq j \leq \min(n, 16)} a_j \widehat{A}_{n-j}. \quad (\text{F.dynamic})$$

The inheritance theorem and induction prove $M_2(n) \leq D_n$, $M_{3/2}(n) \leq D_n^{(3/2)}$, and $A_n \leq \widehat{A}_n$. Each term in the minimum is a valid upper bound; minimizing never relies on mutually compatible extremizers.

Put $\beta_s = 31/20$, $\beta_f = 57/50$. Exact rational comparisons of the finite table give

$$a_{12} < \beta_s^{12}, \quad a_r \leq 2\beta_s^{r-1} \quad (0 \leq r < 12, r \neq 2), \quad a_{12}a_2 < 2\beta_s^{13}, \quad (\text{F.supgates})$$

$$b_{12}^2 < \beta_f^{12}, \quad b_r^2 \leq 4\beta_f^r \quad (0 \leq r < 12). \quad (\text{F.fragates})$$

For $n = m + 1 \geq 3$, write $n = 12s + r$. Unless $r = 2$, (F.supgates) gives $A_n \leq 2\beta_s^{n-1}$. In the exceptional case $s \geq 1$; use $a_{12}^{s-1}(a_{12}a_2)$ instead. Consequently the following ceilings are valid for every $m \geq 2$:

$$\begin{aligned} S_m &= \min\{(8/9)\beta_s^m, (4/9)\widehat{A}_{m+1}\}, & \|H_m\|_\infty &\leq S_m, \\ V_m &= (64/135)D_{m+1} - \mu^2, & \langle (H_m - \mu)^2 \rangle_m &\leq V_m, \\ \mathcal{B}_m &= \frac{3(8/9)^{3/2}}{1 + 2^{3/2}} D_{m+1}^{(3/2)}, & \langle H_m^{3/2} \rangle_m &\leq \mathcal{B}_m. \end{aligned} \quad (\text{F.ceiling})$$

The variance ceiling is positive: $D_n \geq 521/200 > 5/3$ for $3 \leq n < 16$, and $D_n \geq 1737/200$ for $n \geq 16$. The supremum envelope's restriction $m \geq 2$ matters: $\|H_1\|_\infty = 88/63 > (8/9)\beta_s$.

The prefactor in $\mathcal{B}_m$ is less than one: its square is $(512/81)/(9 + 4\sqrt{2}) < 1$. Together with (F.fracgates) and $4\beta_f < 5$, this yields the useful but weaker comparison bound

$$\mathcal{B}_m^2 \leq 5\beta_f^m. \quad (\text{F.fracsimple})$$

The density formula uses the exact table expression (F.ceiling), not (F.fracsimple). Jensen and the exact mean also give $\mathcal{B}_m^2 \geq \mu^3$; this will check the radial counting domain.

For the sixteen-step fractional bound set

$$D_{16,n}^{(3/2)} = b_{16}^{\lfloor n/16 \rfloor} b_{n \bmod 16}, \quad \mathcal{B}_{16,m} = \frac{3(8/9)^{3/2}}{1 + 2^{3/2}} D_{16,m+1}^{(3/2)}, \quad \mathcal{B}_m^{\text{best}} = \min\{\mathcal{B}_m, \mathcal{B}_{16,m}\}. \quad (\text{F.sixteenceiling})$$

Inheritance and the same fan identity prove that each is a valid 3/2-moment ceiling. In particular $\mathcal{B}_m^{\text{best}} \geq \mu^{3/2}$. Both fractional ceilings are retained: their minimum is valid without requiring compatible extremizers.

Use of the moment estimates

The finite maximum, second-moment and direct fractional-moment tables supply the capacity inputs used by the density formulas. The two transfer implementations use distinct recurrences with the same reference law and rational table targets. The product inequality extends these finite bounds to every level. The next section applies them to actual source weights and logarithmic-time convergence.

10. Actual-weight capacity and lower natural density

The parameterized construction supplies integer trajectories reaching one at every fixed clock above the reference ratio, with the same root score (R.score). We now count their distinct starting integers. The aggregate source-weight bound controls weighted sums in residue classes. An elementary maximum estimate already gives a common positive density; the certified moment bounds refine it quantitatively. We count each odd source together with its even multiples, using the actual source weights throughout.

Counting notation

The following lookup collects the quantities already defined in the proof.

Proof stage	Symbols and meanings	Definition
Integer trajectories	w, d, A, C_w, ω_w : inverse word, length, total valuation, affine offset and slope weight	(W.inverse), (W.offset), (W.charge)
Reference distribution	G_q, μ_q, f_q, ρ_q : residue group, probability law, density relative to the uniform measure, and its scaled version	(T.law)
Approximation	ϵ_m, C_6 : uniform error bound from level m to finer levels, and its fixed coefficient	(M.constants), (M.envelope)
Positive root mass	$\Pi, V(N), P$: guaranteed reference survival, total future variation bound, and their positive difference	(R.certificate), (R.residual)
Root family	R_s, W_* : known convergent roots and a lower bound for their reciprocal-root weighted marked mass	(R.pool), (R.score)
Sources and reference factors	a_x, H_m : aggregate weight at integer x , and the weighted reference fan	(N.fan), (N.everyclock)
Moment bounds	$S_m, V_m, \mathcal{B}_m$: upper bounds for the fan maximum, variance, and average 3/2-power	(F.ceiling)

Proof stage	Symbols and meanings	Definition
Fractional refinement	$\mathcal{B}_{16,m}, \mathcal{B}_m^{\text{best}}$: sixteen-step ceiling and the smaller fractional ceiling	(F.sixteenceiling)
Density conversion	$\mathcal{D}(u)$: lower-density bound from an all-scale harmonic mass u	(I.conversion), (D.radial)

The variance bound V_m is distinct from the seed variation bound $V(N)$. In this section U denotes total source weight, not the odd accelerated map, and T denotes a residue-capacity coefficient.

Fix an arbitrary $c > c_0 = 3/\log(4/3)$ throughout this section, and write

$$\mathcal{G}_c = \{x \in \mathbb{N}_{>0} : \tau_1(x) \leq c \log x\}, \quad \mu = 8/9, \quad L = 2 \log 2, \quad \kappa = 16 \log 2/9. \quad (\text{D.constants})$$

The counting density is lower natural density, $\underline{d}(\mathcal{G}_c) = \liminf_{Y \rightarrow \infty} \#(\mathcal{G}_c \cap [1, Y])/Y$. The clock is ordinary unaccelerated Collatz time, and logarithms are natural unless marked $\log_2$.

Harmonic sum bound for the source weights

Fix an odd integer modulus Q , a real $X > 0$, a radius $R > 1$, and a finite set of positive odd integers in $[X, RX)$ carrying weights $0 \leq a_x \leq 1/x$. Put $U = \sum_x a_x$ and $w_r = \sum_{x \equiv r \pmod{Q}} a_x$. Odd representatives of one residue class modulo Q have spacing $2Q$. The monotonicity of $1/t$ and comparison of all but the first summand with their preceding intervals give

$$\sum_{\substack{X \leq x < RX \\ x \text{ odd}, x \equiv r \pmod{Q}}} \frac{1}{x} \leq \frac{1}{X} + \frac{\log R}{2Q}.$$

Equivalently, with

$$T_{X,Q,R} = Q/X + \frac{1}{2} \log R, \quad 0 \leq w_r \leq T_{X,Q,R}/Q, \quad (\text{D.residuecap})$$

every nonnegative residue test ψ satisfies

$$\sum_x a_x \psi(x) \leq T_{X,Q,R} \langle \psi \rangle_Q. \quad (\text{D.harmonic})$$

Here $\langle \cdot \rangle_Q$ is the full average modulo Q , namely $\langle \psi \rangle_Q = Q^{-1} \sum_{r \bmod Q} \psi(r)$. In this local notation the subscript is the modulus itself; elsewhere $\langle \cdot \rangle_m$ averages modulo 3^m . The proof bounds a weighted subset by an ambient sum; it does not infer its empirical distribution from a reference law.

Apply this to the source family produced in [the terminal-source section](#). In addition to (N.everyclock), its construction gives the actual radius $R_X = \mathcal{R}_{b(X)} = 16/(1 - 8^{-b(X)}) \rightarrow 16$, by (N.selector) and (N.scales), uniformly in the common cutoff intervals. The source set is contained in $[X, R_X X)$, before it is enlarged to any prescribed $[X, \Lambda X)$, $\Lambda > 16$. For fixed coarse $m \geq 2$, the terminal level $k = k(X)$ eventually exceeds m , and $3^k/X \rightarrow 0$.

Every F_j in (F.fan) permutes G_k . The mixing estimate (T.mix) for f and the normalization $\rho = (2/3)f$ yield

$$\langle |H_k - H_m \circ \pi_{k,m}| \rangle_k \leq \frac{8}{9} \epsilon_m. \quad (\text{D.fandiscrepancy})$$

Indeed, sum the individual discrepancies after the affine permutations; their factors are $\sum_{j \geq 0} 4^{-j} = 4/3$. Nonnegativity and absolute convergence justify exchanging sums and averages. Harmonic charging at the fine modulus 3^k then proves

$$\mathcal{A}_X \leq \sum_x a_x H_m(x) + \frac{8}{9} T_{X,3^k,R_X} \epsilon_m, \quad T_{X,3^k,R_X} \rightarrow L. \quad (\text{D.coarse})$$

Thus the limiting error bound is exactly $\kappa \epsilon_m$. Fixing $\Lambda > 16$ in the final shell declaration does not replace this error bound by $(4/9) \log \Lambda \epsilon_m$: the physical family still has its sharper radius tending to sixteen. Conversely, one may not assume a literal radius sixteen at finite cutoffs.

Three capacity bounds for the same physical mass

For a fixed level m , let $Q = 3^m$, $T = T_{X,Q,R_X}$. The exact mean of H_m is μ , and its variance is at most V_m from (F.ceilings). Subtract the constant U/Q from each residue mass; the centered marker sums to zero. Then

$$\sum_r (w_r - U/Q)^2 = \sum_r w_r^2 - U^2/Q \leq (TU - U^2)/Q.$$

Cauchy–Schwarz, the maximum bound and weighted Hölder give respectively

$$\begin{aligned} \sum_x a_x H_m(x) &\leq \mu U + \sqrt{V_m U(T - U)}, \\ \sum_x a_x H_m(x) &\leq S_m U, \\ \sum_x a_x H_m(x) &\leq U^{1/3} \left(\sum_x a_x H_m(x)^{3/2} \right)^{2/3} \leq U^{1/3} (T \mathcal{B}_m)^{2/3}. \end{aligned} \tag{D.capacity}$$

The last inequality is (D.harmonic) applied to the nonnegative $3/2$ -power. All three bounds use the same actual source weights. Their validity depends on the source-weight bound (N.charge) established before the residue masses were formed.

Proposition 10.1 (general moment-to-mass conversion). Let $s > 1$ and $B_{s,m} \geq \langle H_m^s \rangle_m$, with $m \geq 2$ fixed. For a source-family score $0 < W \leq 27/2^{27}$, set $p_m = (W - \kappa \epsilon_m)_+$. Then

$$v_{s,m}(W) = \frac{p_m^{s/(s-1)}}{(LB_{s,m})^{1/(s-1)}} \tag{D.generalp}$$

is a limiting lower bound on the available unmarked source mass: every $0 \leq u < v_{s,m}(W)$ is available at every sufficiently large scale, for every fixed containing ratio $\Lambda > 16$.

Proof. At a finite cutoff, weighted Hölder and (D.harmonic) give

$$\sum_r w_r H_m(r) \leq U^{1-1/s} (T_{X,3^m,R_X} B_{s,m})^{1/s}.$$

By (D.coarse), the left side is at least $W - \eta - \kappa \epsilon_m - o(1)$ for a fixed positive graft tolerance η . Since $T_{X,3^m,R_X} \rightarrow L$, inversion gives every strict lower bound in (D.generalp) by first choosing a sufficiently small fixed η , then a sufficiently large cutoff. No positive lower bound is asserted when $p_m = 0$.

Jensen and the mean μ give $B_{s,m} \geq \mu^s > 0$. Since $p_m \leq W < L\mu$,

$$v_{s,m}(W) \leq \frac{p_m}{\mu} \left(\frac{p_m}{L\mu} \right)^{1/(s-1)} \leq \frac{W}{\mu}.$$

This lies in the radial domain (D.domain), by the numerical bounds proved there. At $s = 3/2$, the formula is exactly $p_m^3/(L^2 \mathcal{B}_m^2)$. $\square$

The order is denoted by s here to distinguish it from the residual marked mass p_m . The quantity $B_{s,m}$ is a full-group moment, not an s -norm. Further certified orders may be used in this proposition; the optimized formulas here use the direct $3/2$ tables.

Using complementary residue capacity

Proposition 10.2 (two-sided moment capacity). Let $h_r \geq 0$ on a set of Q residues, with full mean μ and $\langle h^s \rangle \leq B_s$, $s > 1$. Suppose $0 \leq w_r \leq T/Q$ and $\sum_r w_r h_r \geq p$, where $0 < p < T\mu$. Define $q_s(p, T) \in (0, p/\mu]$ by

$$\frac{p^s}{q_s^{s-1}} + \frac{(T\mu - p)^s}{(T - q_s)^{s-1}} = TB_s. \tag{D.twosidedcapacity}$$

There is one such solution on that interval, and $\sum_r w_r \geq q_s(p, T)$. If $B_s = \mu^s$, its value is p/μ .

Proof. Scale the weights down to marked mass exactly p , which can only decrease their total mass; call the result u . Complementary weights $T/Q - w_r$ have total mass $T - u$ and marked mass $T\mu - p$. Hölder on the occupied and complementary parts, followed by addition, gives

$$\frac{p^s}{u^{s-1}} + \frac{(T\mu - p)^s}{(T - u)^{s-1}} \leq TB_s. \tag{D.complementholder}$$

Here $0 < u < T$. On $(0, p/\mu)$ the left side is strictly decreasing: its derivative has the sign of $((T\mu - p)/(T - u))^s - (p/u)^s$. It decreases from infinity to $T\mu^s$, which is at most TB_s by Jensen. This proves the stated solution and lower bound if $u < p/\mu$; if $u \geq p/\mu$, the lower bound is immediate. $\square$

An explicit weaker inverse follows by replacing $(T - u)^{s-1}$ with its upper bound T^{s-1} :

$$\tilde{v}_s(p, T) = \frac{p^{s/(s-1)}}{[T\{B_s - (\mu - p/T)^s\}]^{1/(s-1)}} \leq q_s(p, T). \quad (\text{D.complementexplicit})$$

Its denominator is positive. Indeed, putting $a = p/T \in (0, \mu)$,

$$B_s - (\mu - a)^s \geq \mu^s - (\mu - a)^s \geq a\mu^{s-1} > 0.$$

Consequently $\tilde{v}_s(p, T) \leq p/\mu$. It strictly exceeds the occupied-only bound $p^{s/(s-1)}/(TB_s)^{1/(s-1)}$, since the subtracted complementary term is positive.

At a finite cutoff use $T = T_{X, 3^m, R_X}$, the actual capacity coefficient, and a guaranteed positive marked mass below $W - \eta - \kappa\epsilon_m - o(1)$. First fix a sufficiently small graft tolerance η , then take the cutoff large. Continuity as $T \rightarrow L$ gives every strict lower bound below the displayed inverse with $p = p_m = (W - \kappa\epsilon_m)_+$, $T = L$. Set the inverse to zero when $p_m = 0$. The bound by p_m/μ also gives continuity there and keeps it in (D.domain). Thus the same radial counting applies; no limiting capacity is imposed at a finite cutoff.

Fix a score $0 < W \leq W_{\max} := 27/2^{27}$ for which the local source-family theorem holds. Set

$$p_m = (W - \kappa\epsilon_m)_+, \quad B_m(p) = 2\mu p + V_m L,$$

$$\Psi_m(p) = \frac{2p^2}{B_m(p) + \sqrt{B_m(p)^2 - 4(\mu^2 + V_m)p^2}}, \quad \Psi_m(0) = 0, \quad (\text{D.inverse})$$

$$u_m^{\text{MV}}(W) = \max\{\Psi_m(p_m), p_m/S_m\}, \quad v_m(W) = \frac{p_m^3}{L^2 \mathcal{B}_m^2}, \quad u_m(W) = \max\{u_m^{\text{MV}}(W), v_m(W)\}. \quad (\text{D.massrecipe})$$

The stronger branches and their common maximum are

$$\begin{aligned} v_{16,m}(W) &= \frac{p_m^3}{L^2 \mathcal{B}_{16,m}^2}, \\ \tilde{v}_m(W) &= \frac{p_m^3}{L^2 [\mathcal{B}_m^{\text{best}} - (\mu - p_m/L)^{3/2}]^2} \quad (p_m > 0), \quad \tilde{v}_m(W) = 0 \quad (p_m = 0), \quad (\text{D.refinedmass}) \\ u_{16,m}(W) &= \max\{u_m(W), v_{16,m}(W)\}, \\ u_m^{\text{comp}}(W) &= \max\{u_{16,m}(W), \tilde{v}_m(W)\}. \end{aligned}$$

The complementary branch uses the better of the two moment ceilings. All branches concern the same actual source family and fixed score. Their limiting mass bounds are at most W/μ , as are the original branches by (D.domain).

The label MV denotes the combination of the centered second-moment and maximum bounds; it does not change the root score or mixing estimate. The added term $v_m(W)$ is the direct fractional-moment bound. Each expression estimates the same unmarked source mass. The square root and its choice of branch have the following justification. On $[0, L/2]$, the function $g_m(u) = \mu u + \sqrt{V_m u(L - u)}$ is strictly increasing, with $g_m(0) = 0$ and $g_m(L/2) \geq \mu L/2 > W_{\max}$. Consequently there is a unique inverse for $0 \leq p \leq W_{\max}$. Squaring on the branch $p - \mu u \geq 0$ gives

$$(\mu^2 + V_m)u^2 - (2\mu p + V_m L)u + p^2 = 0.$$

The required solution is the smaller root, whose rationalized expression is (D.inverse). It obeys $\Psi_m(p) \leq p/\mu < L/2$. Choosing the other quadratic root would assert a false capacity bound.

The radial domain will also be respected. Since $S_m \geq \mu$, both benchmark branches are at most W/μ . Jensen in (F.fanmoments) gives $\mathcal{B}_m^2 \geq \mu^3$. Since $\log 2 > 2/3$, $L^2 \mu^3 > (16/9)(512/729) > 1$. Therefore $v_m \leq W^3 \leq W$. In particular

$$0 \leq u_m(W) \leq W/\mu < \frac{1}{2} \log(3/2) < L/2. \quad (\text{D.domain})$$

The strict numerical inequality follows already from $W_{\max}/\mu < 1/100$ and $\frac{1}{2} \log(3/2) > 1/6$, the latter by integrating $1/t > 2/3$ on $[1, 3/2]$.

Lemma 10.3 (all-scale unmarked mass). For each fixed $m \geq 2$, each $u < u_m(W)$ with $u \geq 0$, and each fixed $\Lambda > 16$, every sufficiently large real X has a finite family of odd sources in $\mathcal{G}_c \cap [X, \Lambda X)$, with

$$a_x \leq 1/x, \quad \sum_x a_x \geq u. \quad (\text{D.mass})$$

Proof. Fix m, Λ and a positive graft tolerance η . (N.everyclock) and (D.coarse) give coarse marked mass at least $W - \eta - \kappa\epsilon_m - o(1)$; both fine and coarse capacity coefficients tend to L , uniformly at large cutoffs. Apply (D.capacity). To justify the limit without postulating one zero-error graft, first allow $\eta \downarrow 0$ in these scalar inequalities. If an asserted branch lower bound failed by a fixed positive amount for every sufficiently small tolerance, one could choose $\eta_i \rightarrow 0$ and corresponding sufficiently large bad cutoffs. Their masses have a convergent subsequence, since $U \leq T \rightarrow L$. The limit U_* would satisfy all three inequalities with left side p_m and capacity L . If $U_* \geq L/2$, (D.domain) already contradicts the failure; otherwise monotonicity of g_m and the other two inversions give $U_* \geq u_m(W)$, again a contradiction. Hence a single sufficiently small fixed tolerance and a single associated finite graft work for the chosen strict u and all large X . At $p_m = 0$, all the asserted branch bounds are zero and there is nothing positive to prove. $\square$

The same conclusion (D.mass) holds for every $0 \leq u < u_m^{\text{comp}}(W)$. For the sixteen-step branch use (D.generalp) with (F.sixteen ceiling); for the complementary branch use the finite-cutoff argument after (D.complementexplicit). A finite maximum of valid limiting mass bounds remains valid by choosing an attaining branch and its fixed tolerance. All lie in (D.domain).

These bounds give the explicit density recipes below. [Appendix D.2](#) gives the stronger finite allocation problem using the whole concentration profile, without asserting a new numerical value for the optimized constant.

Joint harmonic allocation and dyadic counting

The *dyadic tower* of an odd integer is the set of its multiples by nonnegative powers of two. We keep these multiples and their source weights together when converting harmonic mass into an integer count.

For positive $x \leq Y$, let

$$D_Y(x) = 1 + \lfloor \log_2(Y/x) \rfloor. \quad (\text{D.dyadic})$$

This is the number of members of that tower at most Y . Distinct odd integers have disjoint dyadic towers. Also $\mathcal{G}_c$ is closed under multiplication by powers of two: for $r \geq 0$, $\tau_1(2^r x) \leq r + \tau_1(x)$, and $c \log 2 > c_0 \log 2 > 1$. Thus all the counted towers remain inside the same literal logarithmic-time set.

Theorem 10.4 (conversion from small harmonic mass to lower density). Suppose, for every fixed $R > 16$ and every sufficiently large real X , the odd members of a dyadically closed set G in $[X, RX)$ admit weights $0 \leq a_x \leq 1/x$ totaling at least u . If $0 \leq u \leq \frac{1}{2} \log(3/2)$, then

$$\underline{d}(G) \geq \mathcal{D}(u) := \frac{32}{225}(e^{2u} - 1). \quad (\text{D.radial})$$

The function $\mathcal{D}$, introduced in (I.conversion), converts weighted counts over intervals of sizes comparable to X into density below a final cutoff. We call it the radial conversion.

Proof. The zero case is immediate. First consider one exact band $(X_j, 16X_j)$, where $X_j = Y/16^j$, $j \geq 1$ is fixed. Away from its four dyadic boundaries, the tower count on $(2^h X_j, 2^{h+1} X_j)$ is $4j - h$, for $h = 0, 1, 2, 3$. Relax integer selection to occupations $0 \leq t_x \leq 1$ at all ambient odd points, taking $t_x = xa_x$ on the actual family and zero elsewhere. Their harmonic mass is $\sum t_x/x$, and their fractional counting cost is $\sum t_x D_Y(x)$. The actual integer count is at least that cost.

The cost per unit harmonic mass is $x D_Y(x)$. In the first subband it increases as $4jx$. The minimum in any later subband is at least $2^h(4j - h)X_j \geq (8j - 2)X_j \geq 6jX_j$. For fixed $1 < s < 3/2$, all odd points in (X_j, sX_j) therefore cost no more per unit mass than any remaining point. An exchange of harmonic mass from a dearer occupied point to a cheaper unoccupied one cannot increase cost. At a prescribed small mass the relaxed minimum is thus obtained by filling the odd points nearest the bottom, with at most one partial occupation, up to the required level.

Uniform integral comparisons and odd spacing give

$$\sum_{\substack{X_j < x < sX_j \\ x \text{ odd}}} \frac{1}{x} = \frac{1}{2} \log s + O(X_j^{-1}), \quad \#\{x \text{ odd} : X_j < x < sX_j\} = \frac{1}{2}(s-1)X_j + O(1). \quad (\text{D.bandcensus})$$

If the prescribed mass tends to a with $0 < a < \frac{1}{2} \log(3/2)$, take $s = e^{2a}$. The $O(X_j^{-1})$ discrepancy in mass changes the minimum cost by only $O(j)$: each available point near this fixed level has mass of order X_j^{-1} and cost of order j . It follows that every occupation of mass at least $a - o(1)$ costs at least

$$2j(e^{2a} - 1) \frac{Y}{16^j} - o(Y). \quad (\text{D.bandcost})$$

At the upper endpoint use every smaller a and let it increase to the endpoint. Removing the finitely many dyadic boundary points costs only $O(X_j^{-1})$ harmonic mass; their extra tower copies cannot weaken a lower count. Thus (D.bandcost) covers all boundary conventions.

The available shells do not have exact radius sixteen. Fix a packing depth $J_0 \geq 1$ and a fixed $R > 16$, then choose the hypothesized families in the pairwise disjoint shells $[Y/R^j, Y/R^{j-1})$, $1 \leq j \leq J_0$. Remove sources below $Y/16^j$. At most

$$\frac{j}{2} \log(R/16) + O(Y^{-1}) \quad (\text{D.stripdebit})$$

of harmonic mass is lost, by the same odd census, with constants allowed to depend on the fixed j, R . The remainder lies in the exact band $[Y/16^j, Y/16^{j-1})$. Its mass is at least $a_j - o(1)$, where $a_j = (u - \frac{j}{2} \log(R/16))_+$. For each positive a_j , apply (D.bandcost) with an arbitrarily small fixed downward tolerance; the zero cases contribute zero. The selected odd sources in different shells are distinct and their towers are disjoint. After dividing their total count by Y , first let $Y \rightarrow \infty$, then $R \downarrow 16$ with J_0 fixed, and finally $J_0 \rightarrow \infty$. All inequalities concern the same set G , so source families and startup cutoffs may depend on these fixed choices. Since $\sum_{j \geq 1} 4j/16^j = 64/225$, the result is

$$\underline{d}(G) \geq \frac{e^{2u} - 1}{2} \sum_{j \geq 1} \frac{4j}{16^j} = \frac{32}{225} (e^{2u} - 1).$$

This proves (D.radial). $\square$

The linear consequence is $\mathcal{D}(u) \geq 64u/225$, which retains the factor $32/15$ over the older separate packing-and-doubling coefficient $2u/15$. The nonlinear expression is stronger still, but its ratio to the linear one tends to one as $u \downarrow 0$. No extra fixed multiplicative improvement is claimed from that small second-order correction.

A common positive density without finite moment certificates

Proposition 10.5 (elementary common-density bound). Suppose (T.mix) has $\epsilon_m \rightarrow 0$, and a fixed persistent pool supplies a score $0 < W \leq W_{\max} = 27/2^{27}$ for (N.everyclock). Choose one fixed $m \geq 2$ such that $\kappa\epsilon_m < W/2$. Then

$$\forall c > c_0 : \quad \underline{d}(\mathcal{G}_c) \geq \frac{4W}{25 \cdot 2^m} > 0. \quad (\text{D.elementarydensity})$$

Both W and m are fixed before c . This deduction uses no moment certificates at level twelve or sixteen.

Proof. Such an m exists because $\epsilon_m \rightarrow 0$. The elementary recursion (B.markercap) and the fan definition (F.fan) give

$$\|H_m\|_\infty \leq \sum_{j \geq 0} 4^{-j} \|\rho_m\|_\infty \leq \frac{8}{9} 2^m. \quad (\text{D.elementarycap})$$

For a fixed $c > c_0$, the source theorem (N.everyclock) and (D.coarse) therefore supply every mass strictly smaller than

$$\frac{W - \kappa\epsilon_m}{(8/9)2^m} > a_m := \frac{9W}{16 \cdot 2^m}.$$

Indeed, for a fixed positive graft tolerance η , the finite coarse marked mass is at least $W - \eta - \kappa\epsilon_m - o(1)$. First make η small enough for the desired strict mass, then take all cutoffs sufficiently large. This works for every fixed containing ratio $\Lambda > 16$, using the actual radius tending to sixteen.

Since $m \geq 2$, $a_m < W \leq W_{\max} < 1/100$; also $\frac{1}{2} \log(3/2) \geq 1/6$, by integrating $1/t \geq 2/3$ on $[1, 3/2]$. Thus every $0 < u < a_m$ is in the radial domain. The set $\mathcal{G}_c$ is dyadically closed as proved above. Apply (D.radial), let $u \uparrow a_m$, and use continuity and $\mathcal{D}(u) \geq 64u/225$:

$$\underline{d}(\mathcal{G}_c) \geq \mathcal{D}(a_m) \geq \frac{64}{225} \frac{9W}{16 \cdot 2^m} = \frac{4W}{25 \cdot 2^m}.$$

The fixed score and coarse level have no dependence on c ; the graft and eventual cutoff may have such a dependence. $\square$

Explicit density bounds

Define

$$c_{\text{MV}}(W) = \max_{m \geq 2} \mathcal{D}(u_m^{\text{MV}}(W)), \quad c_{\text{FM}}(W) = \max_{m \geq 2} \mathcal{D}(u_m(W)). \quad (\text{D.densityrecipe})$$

The two refinements are

$$\begin{aligned} c_{\text{FM},16}(W) &= \max_{m \geq 2} \mathcal{D}(u_{16,m}(W)), \\ c_{\text{comp}}(W) &= \max_{m \geq 2} \mathcal{D}(u_m^{\text{comp}}(W)). \end{aligned} \quad (\text{D.refinedrecipe})$$

Here comp denotes complementary residue capacity.

Here FM stands for fractional moment. The two formulas c_{MV} and c_{FM} optimize the same score, finite tables, mixing error bound and counting function; only the fractional-moment branch is omitted from c_{MV} . These are the complete optimized formulas, rather than the simpler exponential upper bounds used to compare them.

Theorem 10.6 (quantitative density at every larger clock). With the canonical primitive input (M.primitive), survival theorem (R.certificate) and finite certificates (F.finitetarget), every source-family score $0 < W \leq W_{\text{max}}$ from a fixed convergent antichain satisfies

$$\begin{aligned} \underline{d}(\mathcal{G}_c) &\geq c_{\text{comp}}(W) \geq c_{\text{FM},16}(W) \\ &\geq c_{\text{FM}}(W) \geq c_{\text{MV}}(W) > 0 \quad (c > c_0). \end{aligned} \quad (\text{D.headline})$$

In particular the baseline score W_* gives

$$\forall c > c_0 : \quad \underline{d}(\mathcal{G}_c) \geq c_{\text{FM}}(W_*) > 0. \quad (\text{D.everyclock})$$

All scores and density functions are independent of c . The same implications hold for any fixed integer $C \geq 1$ satisfying (M.primitive), with its corresponding mixing, startup and score recipes. For fixed $c > c_0$ and any strict lower bound $d < c_{\text{comp}}(W)$, all sufficiently large X have $\#(\mathcal{G}_c \cap [1, X]) \geq dX$.

Proof. The reference-mixing, seed-mark and persistent-root sections provide the graft premises: (R.graftinput) discharges (G.marks), and (N.everyclock) supplies all-scale source families in the fixed set $\mathcal{G}_c$. For each m , use (D.mass) and its extension to (D.refinedmass), followed by (D.radial). Let the strict mass bound increase to the corresponding limiting value and use continuity of $\mathcal{D}$. Taking the supremum gives each density formula.

Each supremum is a positive attained maximum. The error envelope tends to zero, so $p_m > 0$ eventually. The maximum branch tends to zero by (F.spike), and the centered inverse tends to zero since $V_m \rightarrow \infty$ and $\Psi_m \leq 2W^2/(V_m L)$. Both fractional branches tend to zero because $b_{12}, b_{16} > 1$. Finally $\mathcal{B}_m^{\text{best}} \rightarrow \infty$; once it exceeds $2\mu^{3/2}$,

$$\tilde{v}_m(W) \leq 4 \max\{v_m(W), v_{16,m}(W)\} \rightarrow 0.$$

This follows by bounding the subtracted complementary term by $\mu^{3/2}$. Thus each positive mass sequence tends to zero and attains its supremum at a finite level. Inclusion of branches proves the non-strict comparisons. Section 11 proves the strict improvements.

The pool and score precede c by (N.everyclock); none of L, κ, ϵ_m , the moment ceilings or the optimizations depends on c . Only the graft and eventual cutoffs do. The final counting assertion is the definition of lower density. In particular, $c = 10431/1000$ gives the certified numerical specialization. $\square$

Corollary 10.7 (an unspecified vanishing clock loss). If a constant $d_0 > 0$ satisfies the common-clock conclusion above, there exists a nonincreasing $\eta : [1, \infty) \rightarrow (0, \infty)$, with $\eta(x) \rightarrow 0$, such that

$$\underline{d}\{x \geq 1 : \tau_1(x) \leq (c_0 + \eta(x)) \log x\} \geq d_0. \quad (\text{D.vanishingclock})$$

In particular one set of lower density at least d_0 has $\limsup_{x \rightarrow \infty} \tau_1(x)/\log x \leq c_0$, with the limit restricted to that set. No decay rate for η is specified.

Proof. Put $A_j = \mathcal{G}_{c_0+1/j}$; these sets decrease with j . Choose strictly increasing integer cutoffs $X_j \geq j$ such that $\#(A_j \cap [1, X]) \geq (d_0 - 1/j)X$ for every $X \geq X_j$. Set $\eta = 1/j$ on $[X_j, X_{j+1})$ and $\eta = 1$ below X_1 . If X is in the j -th band, every point of A_j below X satisfies the new clock bound: its earlier band has allowance at least $1/j$. The resulting set therefore has at least $(d_0 - 1/j)X$ points below X . Let $X \rightarrow \infty$. This diagonalization uses the nested fixed-clock sets; the endpoint $c = c_0$ is discussed in Section 5. The same proof applies to any fixed target with a common density constant throughout $c > c_0$. $\square$

Theorem 10.8 (finite-pool transfer to a fixed target). Under the same analytic and moment inputs as (D.everyclock), fix a positive target y_* and a nonempty finite pool satisfying (G.marks), (W.antichain), and the finite root-to-target path hypothesis of (N.targetfamily). If $0 < W \leq W_{\max}$, then

$$\forall c > c_0 : \quad \underline{d}(\mathcal{G}_c(y_*)) \geq c_{\text{comp}}(W) \geq c_{\text{FM}}(W) > 0. \quad (\text{D.finitepool})$$

Without the moment-table premise, the same pool gives the elementary bound (D.elementarydensity), with $\mathcal{G}_c(y_*)$ replacing $\mathcal{G}_c$.

Proof. Apply (N.targetfamily). Every step of the harmonic-capacity and coarse-transfer argument uses its source weights, radius and precision; none depends on which fixed target is eventually visited. The target-clock set is dyadically closed because $\tau_{y_*}(2^r x) \leq r + \tau_{y_*}(x)$ and $c \log 2 > 1$. Thus (D.mass), including its refined branches, the general set formulation of (D.radial), and the attainment argument in (D.headline) give the bound. The elementary maximum argument proves the table-free variant. The pool, score and coarse estimates precede c , so one constant works throughout its stated range. $\square$

Corollary 10.9 (odd sources in every large shell). Fix $c > c_0$, $\Lambda > 16$, $m \geq 2$, and $0 < u < u_m(W)$ for a score in (D.everyclock). Then for all sufficiently large real X ,

$$\#\{x \in \mathcal{G}_c \cap [X, \Lambda X] : x \text{ odd}\} \geq uX. \quad (\text{D.shellcount})$$

Indeed, (D.mass) as used above supplies unmarked mass at least u , and each source has weight at most $1/X$. Marked mass alone would not give this count. The assertion uses $\Lambda > 16$ and a strict mass below the limiting bound.

Every prescribed target not divisible by three

Two complete incoming periods suffice for the general target theorem.

Theorem 10.10 (every fixed positive unit target). For every positive integer y_* , $3 \nmid y_*$, an effectively defined $d_{y_*} > 0$ satisfies

$$\forall c > c_0 : \quad \underline{d}(\mathcal{G}_c(y_*)) \geq d_{y_*}. \quad (\text{D.everytarget})$$

Proof. Use the selected startup N and $q = q_N$, as in (R.markapplication). Write $y_* = 2^v y$, with y odd. In (W.constructor), choose $t_0 \geq 0$ so that $X_{t_0}(y) \geq 16^b$ and $2t_0 + e(y) \geq v + 1$. Take the unit roots in each of the disjoint index intervals $[t_0, t_0 + 3^q)$ and $[t_0 + 3^q, t_0 + 2 \cdot 3^q)$. Each contains one root in every unit residue. All roots have next odd value y . By the cycle-predecessor lemma of Section 1, at most one root in their union is periodic. At least one whole block therefore satisfies (W.antichain).

Let

$$M_{\max} = X_{t_0+2 \cdot 3^q-1}(y), \quad W_{y_*} = \min\{W_{\max}, K_q P / M_{\max}\} > 0, \quad d_{y_*} = c_{\text{FM}}(W_{y_*}). \quad (\text{D.targetscore})$$

The identity (B.physical), uniform over roots, and full future variation bound (R.latermarks) give $z_r = (g_N(r) - V(N))_+$ on either block, with $\sum_r z_r \geq K_q P$. Hence the suitable block has $\sum_M z_M / M \geq K_q P / M_{\max}$. Its root indexed by t reaches y_* along the ordinary odd-and-halving segment in $1 + 2t + e(y) - v$ steps. These specified finite paths need no information about the orbit after y_* . Apply (D.finitepool).

The score recipe is effective and the same for either candidate block; no cycle-detection procedure is needed to define it or to prove that one block is suitable. All choices are finite and precede c . With the qualitative analytic input, (D.elementarydensity) instead gives $4W_{y_*}/(25 \cdot 2^m) > 0$ at one fixed m making the mixing error smaller than half the score, without the large moment tables. $\square$

This proves (I.targets). It does not assert convergence of the target.

Corollary 10.11 (one-period quantitative target score). Retain the selected startup $q = q_N$, $K = K_q$, $B = B_q$, $P = \Pi - V(N)$ and the index t_0 of the preceding proof. Set

$$M_1 = X_{t_0+3^q-1}(y), \quad W_{y_*}^{\text{one}} = \frac{KP - B}{M_1}. \quad (\text{D.oneperiodscore})$$

Then $0 < W_{y_*}^{\text{one}} < 2^{-4096}$, and

$$\forall c > c_0 : \quad \underline{d}(\mathcal{G}_c(y_*)) \geq c_{\text{comp}}(W_{y_*}^{\text{one}}) > 0. \quad (\text{D.oneperioddensity})$$

Its score strictly improves the two-period value in (D.targetscore):

$$\frac{W_{y_*}^{\text{one}}}{W_{y_*}} > \frac{1}{2} 4^{3^q}. \quad (\text{D.oneperiodgain})$$

Proof. Equations (R.denominator) and (R.conductorguards), with $b = 2^{25}$, $N \geq 16$, and $3/2 > \sqrt{2}$, give

$$\frac{KP}{B} = P(3/2)^q > 2^{q/2-1575N-6305} \geq 2^{(b/2-1575)N-6305} > 2. \quad (\text{D.singlecap})$$

The last exponent is at least 268403951. One full unit period has persistent marks totalling at least KP . Remove its periodic root if present. There is at most one, so its removal costs at most B ; the remaining pool is nonempty and satisfies (W.antichain). It retains the specified finite paths to the actual target, including its even factor. Root-uniform persistence gives score at least $(KP - B)/M_1$. Apply (D.finitepool).

For the domain, put $T = 3^q$. The recurrence $X_{t+1}(y) = 4X_t(y) + 1$ gives $M_1 \geq 16^b 4^{T-1}$, while $KP < K < T$ and $T/4^{T-1} < 1$. Hence $W_{y_*}^{\text{one}} < 16^{-b} < 2^{-4096}$. If $M_2 = X_{t_0+2T-1}(y)$, the same recurrence gives

$$M_2 = 4^T M_1 + (4^T - 1)/3 > 4^T M_1.$$

The clipping in (D.targetscore) is inactive by the same domain bound. Since $KP - B > KP/2$, division proves (D.oneperiodgain). The score is a finite rational recipe; no cycle-decision procedure is needed to define it. The two-period construction remains available for the qualitative theorem without the scalar cap comparison. $\square$

Corollary 10.12 (exact obstruction to positive-density hitting). For every positive integer y and every fixed $c > c_0$,

$$\boxed{\underline{d}(\mathcal{G}_c(y)) > 0} \iff 3 \nmid y. \quad (\text{D.targetcriterion})$$

Proof. If $3 \nmid y$, apply (D.everytarget). Suppose $3 \mid y$. An odd predecessor of a number divisible by three is impossible, since $3x + 1 \equiv 1 \pmod{3}$. Its only predecessor is therefore its double, which is again divisible by three. Iterating this observation backwards along any finite path, and conversely taking repeated halvings, gives

$$\begin{aligned} \{x \geq 1 : \tau_y(x) < \infty\} &= \{2^j y : j \in \mathbb{Z}_{\geq 0}\}, \\ \#\{x \leq X : \tau_y(x) < \infty\} &= 1 + \lfloor \log_2(X/y) \rfloor = O_y(\log X) \quad (X \geq y). \end{aligned} \quad (\text{D.divisiblepredecessors})$$

The full predecessor basin has natural density zero, and so does its subset $\mathcal{G}_c(y)$. $\square$

There can be no positive density constant uniform over all admissible targets: the predecessor basins of the infinitely many distinct $R_s > 1$ with $3 \nmid s$ are pairwise disjoint, and these targets are units modulo three. If any J such basins each had lower density at least $d > 0$, their union would have lower density at least $Jd \leq 1$. Letting the finite number J increase rules out a common positive d .

Remark (a score from one persistent root). At least one root M from (R.pool) has a fixed score $W_M > 0$ such that, for every $c > c_0$,

$$\underline{d}\{x : \tau_M(x) \leq c \log x\} \geq c_{\text{FM}}(W_M) > 0, \quad (\text{D.roottarget})$$

where τ_M is the least ordinary hitting time of M .

Proof. By (R.persistence), $\sum_s z_s \geq K_q P > 0$. Choose the first s , in increasing order in the finite set $\mathcal{S}_q$, with $z_s > 0$, put $M = R_s$, and take $W_M = \min\{z_s/M, W_{\max}\} > 0$. These are finite rational operations, although no numerical target is evaluated here. The singleton pool obeys (G.marks) with that score. Every constructed inverse history passes through M , so its bound on time to one also bounds its time to M . The latter logarithmic-time set is dyadically closed. Repeat the proof of (D.everyclock) for this set. This singleton argument selects a positive entry of the persistent vector; it does not assert that every entry is positive. $\square$

Remark (one-block construction for a convergent target). For every fixed convergent odd unit y , the set $\mathcal{G}_c(y)$ has positive lower natural density, with one target-dependent constant for all $c > c_0$. Indeed choose a barrier exceeding 16^b and every odd state on the finite trajectory of y . A complete period of the roots $X_t(y)$ above that barrier is residue-complete and a convergent antichain: after its first odd step every later odd state is below the barrier. The persistent vector has positive total mass on this block. Apply (D.finitepool) with target y .

Corollary 10.13 (qualitative analytic input). A single positive lower density bound for all $c > c_0$ follows from the local primitive theorem and the seed, persistent-root, graft and terminal constructions, without the large finite moment certificates. Alternatively, Tao’s published reference-law estimates [1] suffice without specifying a numerical analytic coefficient.

Proof. For the local route, (E.primitive) supplies (M.primitive), Section 6 gives its mixing envelope, and Sections 7–8 give a fixed positive score. Choose one coarse level at which the error is less than half the score, and apply (D.elementarydensity). This needs no large moment certificate.

For the alternative route, [1, Proposition 1.17] applied at order 6409 gives (M.primitive) for an integer coefficient $C \geq 1$, with the Fourier sign changed by conjugation when necessary. Section 6 and the finite startup argument then give a positive score $W(C)$.

Alternatively, [1, Proposition 1.14] supplies (T.mix) with unhalved error $K_A m^{-A}$. Take $A = 2314/25$ and an integer $\widehat{C} > \max\{K_A, 2^{8838}\}$. The envelope in (B.coefficientscope), the survival bound $\Pi = 24/25$, and the root allocation give a positive score. In either case decrease the score into $(0, W_{\max}]$, choose a coarse level with error less than half that score, and apply (D.elementarydensity).

In every route the coefficient, pool, score and coarse level are fixed before c . The law and norm are those of (T.law) and (T.mix). $\square$

The alternative scores depend on their chosen analytic coefficients. The explicit density formulas use $C_{\text{primitive}}$ and the finite tables of Section 9.

Corollary 10.14 (explicit rational specialization). For rational $0 < W \leq W_{\max}$ satisfying the source-family theorem, let

$$m_{\circ} = \max \left\{ 2, \left\lceil \left(\frac{5C_6}{2W} \right)^{25/2314} \right\rceil \right\}.$$

Then the bound announced in (I.rational) holds:

$$c_{\text{FM}}(W) > \frac{4}{1125} W^3 (57/50)^{-m_{\circ}}. \quad (\text{D.explicit})$$

Proof. The chosen envelope obeys $\epsilon_{m_{\circ}} \leq C_6 m_{\circ}^{-2314/25} \leq 2W/5$. Since $\kappa = 16 \log 2/9 < 100/81 < 5/4$, its residual $p_{m_{\circ}}$ is strictly greater than $W/2$. Using $L^2 < 2$, (F.fracsimple), and the fractional branch in (D.massrecipe), gives

$$u_{m_{\circ}}(W) > \frac{W^3}{80(57/50)^{m_{\circ}}}.$$

The inequality $\mathcal{D}(u) \geq 64u/225$ proves (D.explicit). The logarithmic bounds used here follow from $\log 2 < 25/36$; then $L^2 < (25/18)^2 < 2$. For the canonical integer C_6 and rational W , the least integer in the ceiling is determined by

$$m^{2314}(2W)^{25} \geq (5C_6)^{25}.$$

Clearing denominators makes this an integer comparison. The recipe terminates, although its integers are not of practical size. $\square$

For any constant strictly below the displayed density bound, the proof supplies an eventual all-cutoff lower bound with that constant. A lower density liminf alone does not assert exact attainment of the displayed constant at every sufficiently large finite cutoff. Neither the fixed root pool nor the finite coarse precision is chosen anew at each cutoff.

The estimate combines the source-family theorem, the reference-moment bounds and the counting argument above. Its analytic input is the primitive Fourier theorem (M.primitive); the survival and moment certificates supply the stated finite inequalities. It remains to quantify the improvement over the formula without fractional moments.

11. Quantifying the fractional-moment improvement

We first compare the twelve-step fractional-moment formula with the optimized maximum and second-moment bounds. Its fractional branch has slower growth in the ternary level, but contains an extra power of the small marked mass; the comparison must account for both effects. We then prove the strict improvements from sixteen-step propagation and complementary capacity, retaining the same mixing envelope (M.envelope).

Two properties of the certified error envelope

For every integer $n \geq 2$, the explicit formula (M.envelope) satisfies

$$\epsilon_n \geq \frac{1}{300}n^{-95}, \quad \epsilon_m \leq \frac{n}{m}\epsilon_n \quad (m \geq n \geq 2, \epsilon_n < 2). \quad (\text{Q.error})$$

These inequalities describe the envelope (M.envelope).

To prove them, use $\log 2 < 25/36$. Throughout the compact range $80 \leq v \leq 679/5$, one has $q_v < 95$ and $a(v) > 1$, while $A = 2314/25 < 95$. The latter exponent bound follows already from $6407 - (25/72)(679/5)^2 > 1$. Every nonconstant candidate in (M.envelope) therefore decreases at least as x^{-1} ; in the second term of (M.widthbound), $q_v > 2$ and $1/\log x$ is also decreasing. If $\epsilon_n < 2$, a nonconstant candidate attains that minimum. It remains available at every larger level, including when crossing m_0 , where more choices are added. This proves the second assertion. Compactness and continuity ensure attainment of the width minimum.

For the first assertion, the second positive term of each compact-width candidate is at least

$$\frac{32}{95^2}n^{-95} > \frac{1}{300}n^{-95},$$

using $\log n \leq n$ and $q_v < 95$. The constant two and $C_6 n^{-A}$, with $C_6 \geq 1$, also exceed this floor. Taking the minimum proves (Q.error) without estimating the actual reference discrepancy from below.

A bound for the maximum and second-moment formula

Fix $0 < W \leq 2^{-4096}$, put $a = \log_2(1/W)$, and let

$$n = \min\{m \geq 2 : W > \kappa \epsilon_m\}. \quad (\text{Q.firstpositive})$$

This minimum exists because $\epsilon_m \rightarrow 0$. Since $\kappa > 1$, (Q.error) implies $W > n^{-95}/300$, and hence

$$n > 2^{(a-9)/95} > 1000(2a+32), \quad n > 32. \quad (\text{Q.largeconductor})$$

For the second inequality, at $a = 4096$ the middle expression exceeds $2^{43} > 1000(2 \cdot 4096 + 32)$. The logarithmic derivative of its ratio to $2a + 32$ is $(\log 2)/95 - 1/(a + 16) > 1/190 - 1/4112 > 0$ thereafter. The constant $\kappa > 1$ follows from $\log 2 > 2/3$.

For the collision table, exact rational comparisons give

$$e_{16} > (8/7)^{16}, \quad \frac{64}{135}(7/8)^{14} > \frac{1}{16}.$$

Because every remainder $e_r \geq 1$, writing $m + 1 = 16s + r$, $0 \leq r < 16$, yields

$$(64/135)D_{m+1} \geq \frac{1}{16}(8/7)^m.$$

At $m \geq 32$ this exceeds $2\mu^2$, so

$$V_m \geq \frac{1}{32}(8/7)^m. \quad (\text{Q.variancefloor})$$

This is a lower bound for the ceiling V_m .

If $U = \Psi_m(p_m)$, the correct inverse relation in (D.inverse) gives $p_m = \mu U + \sqrt{V_m U(L - U)}$, with $U < L/2$. Thus $U \leq 2p_m^2/(V_m L) \leq 64W^2(7/8)^m$, since $L > 1$. The other benchmark branch obeys $p_m/S_m \leq (3/2)W(2/3)^m$, by (F.spike), since S_m is a valid upper bound for that maximum. All branches vanish at $m < n$. For every $m \geq n$, both benchmark branches are at most $64W(7/8)^n$. Their actual masses are also less than $1/2$, by (D.domain). On $0 \leq u \leq 1/2$, convexity of the exponential and $e < 3$ give $e^{2u} - 1 \leq 2(e - 1)u < 4u$. Here $e < 3$ follows by summing $1/k! \leq 2^{-(k-1)}$ for $k \geq 2$, with strict inequality at $k = 3$. Therefore

$$c_{\text{MV}}(W) \leq \frac{8192}{225}W(7/8)^n. \quad (\text{Q.oldbound})$$

The comparison bounds the formula c_{MV} .

Strict improvement over the optimized benchmark

Choose $m = \lceil 101n/100 \rceil$, solely for this comparison; the headline still optimizes all levels. The definition of n implies $\epsilon_n < W/\kappa < 2$, so (Q.error) proves

$$p_m = W - \kappa\epsilon_m > W(1 - n/m) \geq W/101. \quad (\text{Q.residual})$$

From (F.fracsimple), $L^2 < 2$, and $\mathcal{D}(u) \geq 64u/225$,

$$c_{\text{FM}}(W) > \frac{64}{2250 \cdot 101^3} W^3 (57/50)^{-m}. \quad (\text{Q.newbound})$$

The bound $L^2 < 2$ follows from $(25/18)^2 < 2$. Dividing by (Q.oldbound), and writing $\beta_f = 57/50$, gives

$$\begin{aligned} \frac{c_{\text{FM}}(W)}{c_{\text{MV}}(W)} &> \frac{W^2}{1280 \cdot 101^3 \beta_f} \left(\frac{8/7}{\beta_f^{101/100}} \right)^n \\ &> 2^{n/1000 - 2a - 31} > 2. \end{aligned} \quad (\text{Q.paid})$$

The middle inequality is exactly the pair of rational inequalities

$$(8/7)^{1000} > 2\beta_f^{1010}, \quad 1280 \cdot 101^3 \beta_f < 2^{31}, \quad (\text{Q.ratechecks})$$

and the last is (Q.largeconductor). Each inequality is verified by integer cross-multiplication in the finite driver. Unlike an interpolation heuristic, the strict exponential separation compensates for the extra W and the fixed displacement of the chosen coarse level.

The auxiliary first positive level can be eliminated from the comparison. For every $0 < W \leq 2^{-4096}$, with $a = \log_2(1/W)$, we obtain

$$\boxed{\frac{c_{\text{FM}}(W)}{c_{\text{MV}}(W)} > 2^{2^{(a-9)/95}/1000 - 2a - 31} > 2.} \quad (\text{Q.explicitgain})$$

Indeed, substitute $n > 2^{(a-9)/95}$ from (Q.largeconductor) into (Q.paid). The other inequality in (Q.largeconductor) gives $2^{(a-9)/95}/1000 - 2a - 31 > 1$. This proves the stated strict factor for the complete density formulas, with no new estimate or choice of roots.

Combining (D.everyclock), (Q.paid) and (R.scoredomain) gives

$$\boxed{\forall c > c_0 : \quad \underline{d}\{x : \tau_1(x) \leq c \log x\} \geq c_{\text{FM}}(W_*) > 2c_{\text{MV}}(W_*) > 0.} \quad (\text{Q.headline})$$

The explicit specialization $c = 10431/1000$ is included. With $c_* = c_{\text{FM}}(W_*)$, (D.everyclock) proves (I.main), while (Q.explicitgain) gives the stronger comparison (I.gain). The bound uses (M.primitive), its specified coefficient, the finite certificates and the auxiliary estimates proved above. Both comparisons keep the root score and mixing estimate fixed.

Improvement from the sixteen-step and complementary bounds

Theorem 11.1 (strict refinement at a fixed score). For $0 < W \leq 2^{-4096}$, with $a = \log_2(1/W)$,

$$c_{\text{comp}}(W) > c_{\text{FM},16}(W) > \frac{1369}{625} c_{\text{FM}}(W), \quad (\text{Q.sixteengain})$$

and

$$\frac{c_{\text{FM},16}(W)}{c_{\text{FM}}(W)} > 2^{2^{(a-9)/95}/30 - 10}. \quad (\text{Q.sixteenscoregain})$$

All formulas have the same root score and mixing envelope.

Proof. Let m_* attain the positive maximum defining $c_{\text{FM}}(W)$. By (Q.paid), its attaining branch is fractional. Write $n = m_* + 1 = 48t + r$, $0 \leq r < 48$. Equation (Q.largeconductor) gives $t \geq 5$. Every fractional remainder is at least one, and all rows are at most b_{16} ; hence

$$D_n^{(3/2)} \geq b_{12}^{4t}, \quad D_{16,n}^{(3/2)} \leq b_{16}^{3t+3}.$$

The exact scalar comparisons are

$$b_{12}^4 > b_{16}^3, \quad 25b_{12}^{20} > 37b_{16}^{18}, \quad b_{12}^{20} > 16b_{16}^{15}, \quad b_{16}^6 < 256. \quad (\text{Q.sixteenguards})$$

Writing $R = b_{12}^4/b_{16}^3$, the first pair implies $D_n^{(3/2)}/D_{16,n}^{(3/2)} \geq R^t/b_{16}^3 > 37/25$. The common residual and fan prefactor cancel, so $v_{16,m_*} > (37/25)^2 v_{m_*}$. Increasing convexity of $\mathcal{D}$, with $\mathcal{D}(0) = 0$, proves the factor $1369/625$ after optimizing.

For the stronger estimate, the last pair gives $R^5 > 16$ and

$$(D_n^{(3/2)}/D_{16,n}^{(3/2)})^2 \geq R^{2t}/b_{16}^6 > 2^{(8/5)t-8} > 2^{n/30-10},$$

since $t \geq (n-47)/48$ and $47/30 + 8 < 10$. Use $n > 2^{(a-9)/95}$ and the same density conversion.

Finally, a maximizing branch of $c_{\text{FM},16}$ is also fractional: $c_{\text{FM},16} \geq c_{\text{FM}} > c_{\text{MV}}$ excludes both nonfractional branches. At that level the positive complementary term makes $\tilde{v}_m$ strictly larger than both occupied-only fractional bounds, since it uses $\mathcal{B}_m^{\text{best}}$. Monotonicity of $\mathcal{D}$ proves $c_{\text{comp}} > c_{\text{FM},16}$. All masses remain in (D.domain). $\square$

How a larger root score improves the full formula

Theorem 11.2 (cubic score comparison). For the fixed retained envelope and finite tables, if $0 < W_1 \leq 2^{-4096}$ and $W_1 < W_2 \leq W_{\text{max}}$, then

$$c_{\text{FM}}(W_2) > (W_2/W_1)^3 c_{\text{FM}}(W_1). \quad (\text{Q.cubicscore})$$

Proof. The positive maximum at W_1 is attained. By (Q.paid), no benchmark branch can attain it, so at a maximizing level m , $c_{\text{FM}}(W_1) = \mathcal{D}(v_m(W_1))$ with $v_m(W_1) > 0$. Put $a = W_2/W_1 > 1$, $\delta = \kappa\epsilon_m$. Then $W_2 - \delta = a(W_1 - \delta) + (a-1)\delta \geq a(W_1 - \delta)$; hence $v_m(W_2) \geq a^3 v_m(W_1)$. Strict convexity, monotonicity and $\mathcal{D}(0) = 0$ give $\mathcal{D}(a^3 u) > a^3 \mathcal{D}(u)$ for $u > 0$. Maximizing at W_2 proves the claim. All intermediate masses are at most $v_m(W_2)$, already in (D.domain). $\square$

The small-score hypothesis is used at the smaller score to prove which branch attains the maximum. Merely including a cubic branch would not suffice. The inequality is a lower comparison between certified formulas; it gives neither their exact ratio nor an upper bound on that ratio. The same argument for a globally active order- s branch gives exponent $s/(s-1)$, provided its dominance is separately proved. For $c_{\text{FM},16}$, such dominance follows from $c_{\text{FM},16} \geq c_{\text{FM}} > c_{\text{MV}}$ at the smaller score. A maximizing branch is one of the two pure fractional branches, whose denominator depends only on its level. The identical proof gives

$$c_{\text{FM},16}(W_2) > (W_2/W_1)^3 c_{\text{FM},16}(W_1). \quad (\text{Q.sixteencubic})$$

The denominator of $\tilde{v}_m$ also depends on the residual; no cubic score comparison for c_{comp} follows from this argument. The pure-fractional comparisons remain available as lower bounds through the retained formulas.

Appendix C supplies the available scores $W_{\mathcal{F}}^{\text{joint}} > \mathfrak{A}_b W_* > E_b W_* > 0$, all below 2^{-4096} . Therefore (D.finitepool), with target one, and (Q.cubicscore) give a strictly stronger common-density bound than the baseline (Q.headline). The explicit comparisons are stated in (C.densitygain). The meaning of W_* is unchanged.

Finite level search

Each density formula can be approximated effectively. Once a positive rational incumbent for the mass maximum has been certified, a finite level cutoff follows from the decreasing tail bound, valid for $m \geq 32$,

$$u_m(W) \leq 64W^2(7/8)^m + \frac{3}{2}W(2/3)^m + 4W^3b_{12}^{-2\lfloor(m+1)/12\rfloor}. \quad (\text{Q.searchtail})$$

The first two terms were proved above without the additional assumption $m \geq n$. For the third, all $b_r \geq 1$ and the squared fan prefactor exceeds $1/4$: indeed $(512/81)/(9+4\sqrt{2}) > 512/1215 > 1/4$. Also $L^2 > 1$. This gives the stated tail, which tends to zero and can be bounded by exact rational arithmetic when W is rational. The actual W_* is rational by (R.greedy) and the finite rational startup recipe. A rational positive incumbent exists, for example by using $\epsilon_m \leq C_6 m^{-A}$ at a sufficiently large explicit m and certified logarithm bounds. Compact-width minimization on a fixed bounded interval uses continuous elementary functions with computable uniform error bounds. The remaining finite maximum can therefore be approximated to any prescribed accuracy, using certified enclosures for its finitely many values.

For the refined formulas, a tail valid for $m \geq 32$ is

$$u_m^{\text{comp}}(W) \leq 64W^2(7/8)^m + \frac{3}{2}W(2/3)^m + 16W^3(b_{12}^{-2\lfloor(m+1)/12\rfloor} + b_{16}^{-2\lfloor(m+1)/16\rfloor}). \quad (\text{Q.refinedtail})$$

Indeed both fractional fan ceilings exceed $2\mu^{3/2}$ there: the common prefactor exceeds $1/2$, every remainder is at least one, and $b_{12}^2, b_{16}^2 > 4$. Thus $\tilde{v}_m \leq 4 \max\{v_m, v_{16,m}\}$. The argument for the third term of (Q.searchtail) applies to each fractional block length. This proves the displayed tail and finite search for both new optimizers. Their expressions are continuous at zero residual by their common bound p_m/μ , so equality decisions at a zero residual are unnecessary for approximation to a requested accuracy.

The absolute scale of the density formulas

Proposition 11.3 (score-to-density scale). Put $a = \log_2(1/W)$ for $0 < W \leq 2^{-4096}$, and define

$$m_+(W) = \max\{2, \lceil (3C_6/W)^{25/2314} \rceil\}.$$

Then

$$\frac{4W}{25 \cdot 2^{m_+(W)}} \leq c_{\text{FM}}(W) \leq 47W \cdot 2^{-2^{(a-9)/95}/6}. \quad (\text{Q.formulascale})$$

Proof. The first positive level n of (Q.firstpositive) satisfies $n > 2^{(a-9)/95} > 32$. In (Q.searchtail), use $b_{12} > 2$, $(8/7)^6 > 2$, $(3/2)^6 > 2$, and $\lfloor (m+1)/12 \rfloor \geq (m-11)/12$. Its three terms are respectively at most $64W^2 2^{-m/6}$, $(3/2)W 2^{-m/6}$, and $16W^3 2^{-m/6}$, since $2^{11/6} < 4$. As $W < 1$, their sum is at most $(163/2)W 2^{-m/6}$. The recipe vanishes below n ; on its radial domain $\mathcal{D}(u) \leq (128/225)u$. Maximizing and using $(128/225)(163/2) = 10432/225 < 47$ proves the upper bound.

At $m = m_+(W)$, (M.envelope) gives $\epsilon_m \leq C_6 m^{-2314/25} \leq W/3$. Since $\kappa < 5/4$, the resulting error is less than $W/2$. The maximum branch has $S_m \leq (8/9)2^m$; its inverse and $\mathcal{D}(u) \geq 64u/225$ give the lower bound, exactly as in (D.elementarydensity). $\square$

These bounds concern the chosen formula. The error-envelope floor used for the upper bound is not a lower bound on true mixing error. For fixed C_6 and tables, as $W \downarrow 0$, they imply

$$\frac{a}{95} - O(1) \leq \log_2 \log_2(1/c_{\text{FM}}(W)) \leq \frac{25a}{2314} + O_{C_6}(1). \quad (\text{Q.logscale})$$

For fixed $H, P > 0$, (R.scorelower)–(R.scoreupper) give $W(q, P) = 2^{-\Theta(3^q)}$ as $q \rightarrow \infty$. Multiplying by the fixed $\mathfrak{A}_b$ retains this score scale; composition with (Q.formulascale) gives $c_{\text{FM}}(\mathfrak{A}_b W(q, P)) = 2^{-2^{\Theta(3^q)}}$. This describes the formulas at fixed parameters as the residue level grows. The canonical level satisfies the domain bound (C.domainguard).

The upper bound in (Q.formulascale) retains its literal twelve-step meaning. For the strongest retained formula, a separate estimate is

$$c_{\text{comp}}(W) < 55W \cdot 2^{-\frac{1}{8} 2^{(a-9)/95}} \quad (0 < W \leq 2^{-4096}, a = \log_2(1/W)). \quad (\text{Q.refinedscale})$$

To verify it, the rational rows satisfy $b_{12}^2 > (11/10)^{12}$, $b_{16}^2 > (11/10)^{16}$, and $b_{12}^2, b_{16}^2 < 6$. For either block length $d = 12, 16$,

$$b_d^{-2\lfloor (m+1)/d \rfloor} \leq b_d^2 (10/11)^{m+1} < 6(10/11)^m.$$

Hence each pure fractional mass is at most $24W^3(10/11)^m$, and the complementary one at most $96W^3(10/11)^m$ for $m \geq 32$. The two nonfractional branches are bounded by $64W^2(7/8)^m$ and $(3/2)W(2/3)^m$, so their maximum together with the complementary branch is at most $96W(10/11)^m$, since $W < 1$. All branches vanish below the first positive level n . Use $\mathcal{D}(u) \leq (128/225)u$, $(128/225)96 < 55$, $(11/10)^8 > 2$, and $n > 2^{(a-9)/95}$. This proves (Q.refinedscale). The lower bound in (Q.formulascale) remains valid by inclusion of the old formula. Thus the two-sided logarithmic scale in (Q.logscale) also holds for $c_{\text{FM},16}$ and c_{comp} , with their own constants. These upper bounds concern the defined density formulas.

Comparison with Mazur's printed bound

With $L_3(x) = \log_2 \log_2 \log_2 x$, (E.coefficientcomparison) gives $L_3(C_{\text{primitive}}) < 176$ and $L_3(C_{\text{primitive}}^{\text{Maz}}) > 2^{172}$. [Appendix B](#) proves that this coefficient decrease strictly improves the baseline density formula, accounting for both startup rounding and score allocation.

We compare the actual arithmetic recipes, retaining the root parameters printed in [3]. The printed primitive coefficient remains $C_{\text{primitive}}^{\text{Maz}}$, while ours is the smaller local coefficient in (M.canonicalC). The printed sixth-order adapter is C_{Maz} from (M.sourceadapter), while our proved adapter is $C_6 = 2D + 2 < C_{\text{Maz}}$. To preserve the literal printed benchmark, write

$$F_{b,\text{Maz}} = 2^{467} b 16^b (C_{\text{Maz}} + 1).$$

The marker normalization contributes the single factor $2/3$ in (M.marker); it is not absorbed into either coefficient.

Write $b_{\text{Maz}} = 2^{80}$, $H_{\text{Maz}} = 2b_{\text{Maz}} + 1$, and

$$t_{\text{Maz}} = \lceil \log_2 F_{b_{\text{Maz}}, \text{Maz}} \rceil, \quad N_{\text{Maz}} = 20000(t_{\text{Maz}} + 64),$$

Starting from $\beta_0 = b_{\text{Maz}}$, let $\beta_{j+1} = \beta_j + \lfloor \beta_j/100 \rfloor$ and

$$q_{\text{Maz}} = \sum_{j < N_{\text{Maz}}} (\beta_j + \lfloor 3\beta_j/5 \rfloor) + \lfloor \beta_{N_{\text{Maz}}}/4 \rfloor.$$

These expand the recurrences in [3, equations (8.1)–(8.4)]. Its printed constant in equations (8.5)–(8.6) is

$$\mathcal{M} = 4^{H_{\text{Maz}} + 3q_{\text{Maz}}}, \quad m_{\text{Maz}} = \lfloor (88\mathcal{M}C_{\text{Maz}})^{1/6} \rfloor + 1, \quad c_{\text{Maz}} = \frac{3}{256\mathcal{M}2^{m_{\text{Maz}}}}. \quad (\text{Q.mazrecipe})$$

Proposition 11.4 (comparison with the printed constant). The selected score satisfies

$$W_* > \mathcal{M}^{-1}, \quad c_{\text{MV}}(W_*) > \frac{22208}{825} c_{\text{Maz}}, \quad c_{\text{FM}}(W_*) > \frac{44416}{825} c_{\text{Maz}}. \quad (\text{Q.mazgain})$$

Proof. Our construction starts at $b = 2^{25}$, with $t = \lceil \log_2 F_b \rceil$, $N \leq 4t/5$, $q = q_N$, and $H = 2b + 1$. Since $F_{b, \text{Maz}} > F_b$ and these formulas increase with b , $t_{\text{Maz}} \geq t$ and $N_{\text{Maz}} > N$. The floor recurrence, the summands in (B.conductor), and the terminal quarter are nondecreasing in their initial value. The extra positive summand gives $q_{\text{Maz}} \geq q + 1$; also $H_{\text{Maz}} > H$.

By (R.scorelower), (R.denominator), and (R.conductorguards),

$$\begin{aligned} \log_2(W_*\mathcal{M}) &> 3 - 1575N - 6305 + q \log_2 3 + 2(H_{\text{Maz}} - H) + 2(3^{q_{\text{Maz}}} - 3^q) \\ &> 4 \cdot 3^q - q - 6302 > 0. \end{aligned}$$

Here $q \geq bN \geq 1575N$. The final expression is already positive at nine and increases thereafter, whereas our levels exceed nine. This proves the first assertion.

Set $m = m_{\text{Maz}} \geq 2$. The chosen envelope (M.envelope), with exponent $2314/25 > 6$, gives

$$\epsilon_m \leq C_6 m^{-6} < C_{\text{Maz}} m^{-6} < \frac{1}{88\mathcal{M}}, \quad p_m > \frac{347}{352\mathcal{M}},$$

using $\kappa < 5/4$ and $W_* > \mathcal{M}^{-1}$. Since $S_m \leq (8/9)2^m$ and $\mathcal{D}(u) \geq 64u/225$,

$$c_{\text{MV}}(W_*) \geq \mathcal{D}(p_m/S_m) \geq \frac{8}{25} p_m 2^{-m} > \frac{347}{1100\mathcal{M}2^m} = \frac{22208}{825} c_{\text{Maz}}.$$

The radial domain is (D.domain). Finally (Q.paid), applied to W_* using (R.scoredomain), gives the last assertion. $\square$

Retaining the symbolic score in the same calculation gives a stronger comparison. Put $R = W_*\mathcal{M}$ and $E_q = 4 \cdot 3^q - q - 6302$. The proof above gives $R > 2^{E_q} > 1$ and, at $m = m_{\text{Maz}}$, $p_m > (R - 5/352)/\mathcal{M}$. Therefore

$$\frac{c_{\text{MV}}(W_*)}{c_{\text{Maz}}} > \frac{2048}{75} \left(R - \frac{5}{352} \right), \quad \frac{c_{\text{FM}}(W_*)}{c_{\text{Maz}}} > \frac{4096}{75} \left(R - \frac{5}{352} \right). \quad (\text{Q.mazsymbolic})$$

The first bound uses $(8/25)/(3/256) = 2048/75$; the second uses the certified factor greater than two at the same W_* . Substituting 2^{E_q} for R preserves strict lower bounds. The comparison is symbolic in the score and conductor.

Thus the conservative factor over the printed constant is $44416/825 \approx 53.8376$. Against the benchmark obtained by doubling that constant through dyadic closure, the factor is $22208/825 \approx 26.9188$. The explicit numerical clock specialization here is $10431/1000 = 10.431$, compared with $523/50 = 10.46$ in [3]; both count ordinary Collatz steps and use natural logarithms.

This printed benchmark is distinct from $c_{\text{MV}}(W)$, which fixes our root family and optimizes the maximum and second-moment estimates. The former comparison measures the change from the stated earlier bound; the latter isolates the direct fractional-moment gain. The benchmark [3] also includes a single explicit finite-cutoff recipe. Here the proof gives effective eventual thresholds for every strictly smaller density constant, without assembling them into a comparable final cutoff formula.

Further quantitative improvements

The convex-order inequality (F.convexfan) and optimal allocation formula (D.hinge) in [Appendix D](#) identify a further route: certify upper-tail profiles, then compare the resulting density bound with the complete $c_{\text{comp}}(W_*)$, keeping the roots and mixing envelope fixed. The challenge is to propagate the finite distributions efficiently while preserving certified upper bounds and to control the remaining level search. A quantitative comparison can use the exact two-sided inverse (D.twosidedcapacity) as well as propagated profile bounds.

Other direct moment orders can be used in (D.generalp). Interpolating existing moment ceilings is insufficient on its own. Indeed, suppose $1 < s_0 < s_1$, $0 \leq \alpha \leq 1$, $s = \alpha s_0 + (1 - \alpha)s_1$, and the only new bound is $B_s = B_{s_0}^\alpha B_{s_1}^{1-\alpha}$. For a positive residual p , let $v_s = p^{s/(s-1)}/(LB_s)^{1/(s-1)}$. Its logarithm is the weighted average of $\log v_{s_0}$ and $\log v_{s_1}$, with weights $\alpha(s_0 - 1)/(s - 1)$ and $(1 - \alpha)(s_1 - 1)/(s - 1)$. Thus it cannot exceed the better endpoint branch. New direct certificates or additional distribution information are needed. The rate statement (F.divergencerate) does not by itself improve the capacity formula as the order decreases. For exact fan moments and residual $p > 0$, its full logarithm is

$$\log v_s = \frac{s}{s-1} \log(p/\mu) - \frac{\log L}{s-1} - D_s(H_m/\mu),$$

where $D_s(h) = \log \langle h^s \rangle_m / (s - 1)$ for a mean-one density h . The residual and capacity terms remain in this formula alongside the divergence. Adding a certified branch to the maximum can never decrease the bound; a strict gain still needs proof against the whole optimized formula. Polynomial or subexponential collision and fractional-moment bounds remain open.

Absolute density also depends on the root score. From (D.domain),

$$c_{\text{FM}}(W) \leq \frac{32}{225} (e^{9W/4} - 1) \sim \frac{8}{25} W \quad (W \downarrow 0). \quad (\text{Q.scoreceiling})$$

This bounds the certified formula, not the true convergence density. Even the exact residue-capacity optimum satisfies $\mathcal{C}_m(p_m) \leq W/\mu$, by the feasible uniform allocation in (D.hinge). Thus a substantial absolute lower bound requires improving the root score or using further constraints on the actual sources, in addition to improving concentration estimates.

A separate analytic route concerns the head localization in Section 6. With the present squared-width loss, its two competing exponents are

$$a(v) = 6407 - \frac{1}{2}(\log 2)v^2, \quad q_v - 1 = (\log 2)v - 1. \quad (\text{Q.headtradeoff})$$

The first decreases and the second increases with $v > 0$. Their equality determines the best asymptotic exponent within this specific tradeoff; $v_* = 6749/50$ gives the retained exponent $2314/25$ with explicit rational slack. Improving the quadratic localization loss, or replacing that argument by a multiscale or exponential-moment estimate, could improve the resulting exponent. Such a change must reprove the head-frequency estimate, rejection bound, level induction and complete envelope with explicit constants.

The parameterized clock is established in (N.everyclock)–(D.everyclock), with the full fixed density bound for every $c > 3/\log(4/3)$. The endpoint clock remains unproved. A family of bounds valid for every strictly larger coefficient does not imply the endpoint, and the reference-model mean does not establish deterministic optimality below it.

Appendix A. Comparison of fixed-depth and first-crossing seeds

This appendix compares the fixed-depth seed with a central first-crossing benchmark. The constants of both constructions are specified below. Throughout this appendix $b = 2^{25}$, the sequence b_j and the root index origin $H = 2b + 1$ are fixed. The following comparison is uniform in a fixed integer primitive coefficient $C \geq 1$. Put $D = C20^{6409}$; use the same C in both seed recipes. The canonical instance takes $C = C_{\text{primitive}}$. The comparison with [3] takes $C = C_{\text{primitive}}^{\text{Maz}}$. Each comparison uses one common value of C . In this appendix W_* denotes the fixed-depth score at that coefficient; it agrees with (R.score) in the canonical instance.

The two seed recipes

Use $+$ for the fixed-depth seed and its local adapter, and $-$ for the first-crossing benchmark:

$$C_6^+ = 2D + 2, \quad C_6^- = 2D + 2 + 2^{481} =: C_{\text{fc}}(C),$$

$$q^+(n) = \sum_{j < n} b_j + \lfloor b_n/4 \rfloor, \quad q^-(n) = \sum_{j < n} (b_j + \lfloor 3b_j/5 \rfloor) + \lfloor b_n/4 \rfloor. \quad (\text{Q.seedrecipies})$$

The level q_n in (B.conductor) is $q^+(n)$. The first-crossing survival certificate uses $\Pi^- = C_{\text{fc},\text{surv}}/2^{1024}$, with the literal integer $C_{\text{fc},\text{surv}}$:

```
20617628314526275611176215079302618796161541826143958813539212954061565810783736808553005
↪ 0940902298801554929847379180509254740126267558628283858415578260470632821917045223783
↪ 42863853189519310456714109487194639513012960066163451504920988540481985
```

Direct integer comparison gives

$$3 \cdot 2^{-214} < \Pi^- < \frac{1}{4} < \Pi^+ := \frac{24}{25}. \quad (\text{Q.oldsurvival})$$

This integer belongs to the first-crossing survival certificate. The fixed-depth survival proof (R.certificate) does not use that computation.

Set $F_b^\pm = 2^{467}b16^b(C_6^\pm + 1)$ and $t^\pm = \lceil \log_2 F_b^\pm \rceil$. These rounded exponents agree exactly. The common C is an integer, so $2D$ is divisible by 2^{12819} . Both $2D + 3$ and $2D + 3 + 2^{481}$ lie strictly between $2D$ and $2D + 2^{12819}$. There is no power of two in this interval: any such power would exceed $2D \geq 2^{12819}$ and would itself be divisible by 2^{12819} . Multiplication by $2^{467}b16^b$, a power of two at $b = 2^{25}$, preserves the equality of the ceiling logarithms. Hence

$$t^+ = t^- = t, \quad V^+(n) = V^-(n) = V(n), \quad L_t = \lceil \log_2(t + 1) \rceil. \quad (\text{Q.roundinginvariance})$$

The retained conservative variation majorant is the same for this comparison; its separate optimization has not been used.

Define the first-crossing startup and scores by

$$N^- = \min\{n \geq 16 : V(n) < \Pi^-\}, \quad P^- = \Pi^- - V(N^-), \quad W_{\text{fc}} = W(q^-(N^-), P^-),$$

$$N_{\text{fixed}}^- = \lceil 100(t + 3L_t + 214)/131 \rceil = N_{\text{fixed}}, \quad W_{\text{fixed},\text{fc}} = W(q^-(N_{\text{fixed}}^-), \Pi^- - 2^{-213}). \quad (\text{Q.previoussscore})$$

These scores use the first-crossing level and survival constant; $W_{\text{fixed},\text{fc}}$ also fixes the residual budget in advance. Both scores are positive: (R.searchendpoint)'s first three bounds put $V(N_{\text{fixed}}^- - 1) < 3 \cdot 2^{-214} < \Pi^-$. The same term-ratio argument also gives

$$V(N_{\text{fixed}}^-) < \frac{1}{2}V(N_{\text{fixed}}^- - 1) < \frac{32}{25}2^{-214} < 2^{-213}. \quad (\text{Q.previousadmissible})$$

Thus the fixed-budget allocation is admissible.

Level reduction and earlier startup

At every $n \geq 16$,

$$q^+(n) < \frac{63}{100}q^-(n). \quad (\text{Q.conductortrim})$$

To prove it, set $S = \sum_{j < n} b_j$ and $a_n = \lfloor b_n/4 \rfloor$. The recurrence gives $b_n \leq b + S/100$, while $S \geq nb$, so $a_n/S \leq 1/(4n) + 1/400 \leq 29/1600$. Also

$$\sum_{j < n} \lfloor 3b_j/5 \rfloor \geq \left(\frac{3}{5} - \frac{4}{5b}\right)S.$$

Since $(1+x)/(1+\delta+x)$ increases in x for $\delta > 0$,

$$\frac{q^+(n)}{q^-(n)} \leq \frac{1 + 29/1600}{8/5 - 4/(5b) + 29/1600} < \frac{63}{100};$$

the last inequality is rational at $b = 2^{25}$. This is a reduction greater than 37 percent at a common stage. Moreover $b_n = b + S/100 + O(n)$ and $n = o(b_n)$, by geometric growth, so the same-stage ratio tends to 401/641.

Write $N^+ = N$ for the fixed-depth startup. The tail at 16 exceeds one after multiplication by 2^t , so $N^- \geq 17$. For $n \geq 1$,

$$T_3(n-1, \lambda) = \lambda^{-1} \sum_{k \geq n} k^3 \lambda^k < \lambda^{-1} T_3(n, \lambda).$$

Together with $\lambda^{-1} < 5/2$ and (Q.oldsurvival), this gives

$$V(N^- - 1) < (5/2)\Pi^- < 5/8 < 24/25, \quad N^+ < N^-. \quad (\text{Q.earlierseed})$$

In particular $q^+(N^+) < (63/100)q^-(N^-)$ as well.

The fixed-depth startup bound (R.startupoptimum), evaluated at $n = N^-$ and $p = P^-$, followed by (R.conductormonotone), gives

$$W_* > 2^{12193}W(q^+(N^-), P^-) > 2^{12193}W_{\text{fc}}. \quad (\text{Q.seedscoregain})$$

The corresponding fixed-budget comparison follows from the same allocation estimates. In the following denominator, a, d are the rate integers of (R.variationdata). The positive rational $\Pi^- - V(N^-)$ has denominator dividing $2^{1024}d^{N^-}(d-a)^4$, and so exceeds $2^{-1575N^- - 7324}$. The bounds $q^-(n) \geq bn$, $H \leq 8q^-(n) + 7$ and $q^-(n+1) \geq q^-(n) + 1$ also hold. Apply the proof of (R.startupoptimum) with q^- , N^- and that denominator bound. Since $N^- < N_{\text{fixed}}^-$, it proves

$$W_{\text{fc}} > 2^{12193}W_{\text{fixed,fc}}, \quad W_* > 2^{24386}W_{\text{fixed,fc}}. \quad (\text{Q.historicalscoregain})$$

Passing the gains to the complete density formulas

Here $L = 2 \log 2$ is the density interval length, distinct from L_t . For fixed m , the function used in (D.inverse) is increasing and concave on $[0, L/2]$, and its inverse is increasing and convex. Consequently the maximum branch, the centered second-moment branch and the direct fractional-moment branch are convex in W ; the last is a positive multiple of a cubed positive part. Their maximum, its composition with the increasing convex $\mathcal{D}$, and the supremum over levels remain convex. Both density formulas vanish at zero. For either formula c ,

$$c(aW) \geq a c(W) \quad (a \geq 1, 0 \leq W \leq aW \leq W_{\text{max}}). \quad (\text{Q.scoregain})$$

Let $c_{\text{FM},\pm}$ and $c_{\text{MV},\pm}$ denote the complete formulas with adapter $C_6^\pm$; all reference tables and the compact-width terms E_v are identical. Thus $\epsilon_m^+ \leq \epsilon_m^-$ and each $+$ formula dominates its $-$ counterpart at fixed score. Within this comparison, formulas without subscripts are the $+$ formulas at the chosen common C ; at the canonical instance these are the main text's local formulas.

All scores above lie in $(0, 2^{-4096})$, by (R.scoreupper) and their large levels. The quantitative comparison (Q.paid) holds for either adapter: its proof uses the same compact envelope and a coefficient at least one. Applying (Q.scoregain) with the actual score ratios, (Q.seedscoregain) and (Q.historicalscoregain) now gives

$$c_{\text{FM}}(W_*) > 2^{12193}c_{\text{FM},-}(W_{\text{fc}}), \quad c_{\text{FM}}(W_*) > 2^{24387}c_{\text{MV},-}(W_{\text{fixed,fc}}) > 0. \quad (\text{Q.retainedgain})$$

The second inequality includes the certified fractional-moment factor greater than two. These inequalities hold separately at each common coefficient C . For the coefficient in [3], use $C = C_{\text{primitive}}^{\text{Maz}}$ throughout this appendix, then apply (D.coefficientgain) to its left-hand current-seed formula. Thus the main text's local formula also retains the displayed strict gains over those specified first-crossing bounds. The rounding identity (Q.roundinginvariance) is used at a common primitive coefficient.

Appendix B. Dependence on the analytic coefficient

Fix the reference law, seed schedule, and finite moment tables. This appendix compares the resulting startup, root score, and density bound as the valid primitive-decay coefficient varies.

Dependence on the analytic coefficient

Write $t(C), V_C, N(C), P(C), W(C)$ for the fixed-depth startup and allocation of Section 8 with integer primitive coefficient $C \geq 1$ and adapter $2C20^{6409} + 2$. The seed, its survival probability, the rate λ , and the level sequence $q(n)$ are fixed. The coefficient guards in Sections 6–8 hold for every such C . If $C_1 \leq C_2$, then

$$t(C_1) \leq t(C_2), \quad V_{C_1}(n) \leq V_{C_2}(n), \quad N(C_1) \leq N(C_2), \quad W(C_1) \geq W(C_2). \quad (\text{R.coefficientmonotonicity})$$

Moreover $t(C_1) < t(C_2)$ implies $W(C_1) > W(C_2)$.

Proof. The first three comparisons follow from the increasing adapter, the ceiling logarithm and the definition of the first positive residual. If the startup indices agree, the smaller variation tail gives a larger residual at that same level; use the strict increase in (R.greedy). If $N(C_1) < N(C_2)$, apply (R.startupoptimum) for C_1 to

the later level $N(C_2)$ and budget $P(C_2) \leq 1$. Its proof is uniform in the integer coefficient: (R.denominator) is unchanged because $2^{t(C)}$ is an integer. It gives even $W(C_1) > 2^{12193}W(C_2)$. If the rounded exponents differ but startup indices agree, the residual comparison is strict. $\square$

The arithmetic bounds (E.coefficientcomparison) give $C_{\text{primitive}}^{\text{Maz}} > 4C_{\text{primitive}}$. To see that the factor four is allowed, (E.local-size) implies $L_3(4C_{\text{primitive}}) < 177 < 2^{172}$. Consequently the ratio of the corresponding F_b 's exceeds two, so $t(C_{\text{primitive}}^{\text{Maz}}) \geq t(C_{\text{primitive}}) + 1$. Thus

$$W_* = W(C_{\text{primitive}}) > W(C_{\text{primitive}}^{\text{Maz}}) > 0. \quad (\text{R.localcoefficientgain})$$

The comparison includes both the rounded startup and the resulting allocation residual.

Density bounds

Corollary B.1 (coefficient improvement through the complete formula). For the same reference tables, let $c_{\text{FM},C}$ denote the density formula with coefficient C . If $C_1 \leq C_2$ are valid integer primitive coefficients, every term of the mixing envelope for C_1 is no larger than the corresponding term for C_2 . The maximum and moment branches therefore give

$$c_{\text{FM},C_1}(W) \geq c_{\text{FM},C_2}(W) \quad (0 \leq W \leq W_{\max}).$$

The score scaling (Q.scoregain), whose proof uses only the defining functions, then implies, with $W_i = W(C_i)$,

$$c_{\text{FM},C_1}(W_1) \geq \frac{W_1}{W_2} c_{\text{FM},C_2}(W_2). \quad (\text{D.coefficientgain})$$

Both scores are in the required small-score domain by (R.scoredomain). For $C_1 = C_{\text{primitive}}$, $C_2 = C_{\text{primitive}}^{\text{Maz}}$, (R.localcoefficientgain) makes this a strict improvement of the baseline density formula. The larger coefficient is also a valid primitive bound by (E.primitive) and (E.coefficientcomparison). The other retained branches remain valid with the new envelope, and every fixed-target source construction still fixes its new startup, pool and score before choosing the clock.

Appendix C. Proofs and certificates for the finite-root refinement

This appendix proves the finite-root refinement (I.rootdensity). The finite-pool theorem (D.finitepool) and the allocation comparison (R.blockcomparison) reduce that task to residue coverage, separation of root trajectories, and a lower bound for the combined score. We verify these properties for valuation redistributions of certified trajectories and evaluate their parameter counts exactly.

Section C.1 fixes the notation. Sections C.2–C.5 supply the constructions and counting lemmas; Section C.6 assembles the score comparison. Section C.7 gives the canonical finite certificates. The final two sections explain a strict shared-allocation gain and the limitations of increasing the number of blocks at fixed height. The baseline score W_* and the density formulas in the main text retain their definitions.

C.1. Setup and notation

In this appendix, a *seed* is a certified finite trajectory used to construct additional roots. It is distinct from the fixed-depth word family used for the initial marks in Sections 4 and 7.

The constructor (W.constructor) supplies each incoming block. We redistribute valuations along certified finite trajectories while keeping complete residue coverage and a finite root-to-one clock. Throughout this appendix, $b \geq 2^{25}$, $H = 2b + 1$, and the residue period and unit index set are those of the root-pool section. All blocks use the same analytic inputs.

Let a seed have a certified first-hitting trajectory

$$y_0 \xrightarrow{a_1} y_1 \xrightarrow{a_2} \dots \xrightarrow{a_D} 1, \quad 3y_{i-1} + 1 = 2^{a_i}y_i,$$

and let $A_j = \sum_{i \leq j} a_i$, $A_0 = 0$. Fix $e = e(y_0)$, $d \geq 0$, and

$$2^e y_0 \equiv 4^d \pmod{9}, \quad \alpha_0 = 2^e y_0 / 4^d > 3/4. \quad (\text{C.seed})$$

The finite base-height condition and the certified trajectory-height condition must also hold:

$$(2^e y_0 4^{H-d} - 1)/3 \geq 16^b, \quad \max_j y_j < 16^b, \quad d \leq H. \quad (\text{C.baseheight})$$

The base conditions imply $T_0 = 2b - d + e - 1 \geq 0$. Its only possible negative case is $e = 1, d = H$, when the base root $(2y_0 - 1)/3 < y_0 < 16^b$, contradicting the height condition.

For a prefix of length r , put

$$L_j^{(r)} = \frac{2^{A_r - A_j} y_r}{3^{r-j}}, \quad \beta_r = \frac{2^e}{4^d} L_0^{(r)}, \quad \Gamma_r = \frac{2^e}{4^d} \max_{0 \leq j < r} L_j^{(r)}. \quad (\text{C.leading})$$

The depth D and prefix length r below refer only to the certified seed trajectory. They are unrelated to the reference moments and divergence notation used elsewhere.

Notation	Role in this appendix
$y_0, \dots, y_D = 1; a_i, A_j$	Certified seed trajectory, its valuations and their partial sums.
$b, H = 2b + 1; e, d$	Fixed height parameter, initial root index, and seed-dependent parameters in (C.seed).
r, ℓ_i, S_j	Prefix length, added valuations and their partial sums in the general construction below.
Y_j, M_s	Modified trajectory states and the incoming root with index s .
$L_j^{(r)}, \beta_r, \Gamma_r$	Leading height coefficients defined in (C.leading); $L_j^{(r)}$ is unrelated to the counting constant L .
χ_r, T_r	Height correction and remaining increment budget, defined in (C.heightdebit) and (C.budget).
$N_r, Q_r, \Lambda_{y,r}$	Rectangular count, quotient-simplex count and selected layer count, defined in (C.rectcount), (C.quotientcount) and (C.selectedcount).
$E_b, \mathfrak{A}_b$	Rectangular benchmark and canonical guaranteed score multiplier.

C.2. Rectangular redistribution and retained layers

The rectangular rule takes

$$t_i = 3^{i+1} k_i, \quad h = \sum_i t_i \leq T_0 := 2b - d + e - 1,$$

$$Y_r = y_r, \quad Y_{i-1} = \frac{2^{a_i + 2t_i} Y_i - 1}{3}, \quad M_s = \frac{2^{2(s-d-h)+e} Y_0 - 1}{3}. \quad (\text{C.rectangular})$$

A sufficient rectangular condition is $\Gamma_r < 1$. It implies $\beta_r < 1$.

Proof of its arithmetic interface. The valuation identity $\nu_3(4^t - 1) = 1 + \nu_3(t)$ gives $4^{t_i} \equiv 1 \pmod{3^{i+2}}$. Backwards induction gives $Y_j \equiv y_j \pmod{3^{j+2}}$. In particular $Y_0 \equiv y_0 \pmod{9}$, and h is a multiple of nine, so the root's compensated affine coefficient is one modulo nine. The original unit index set is therefore unchanged.

The geometric and clock arguments are the special case $\ell_i = 2t_i$ of the height and clock lemmas below. Every new root is at least 16^b , all later odd states are below 16^b , and its clock is $2s + e + 1 + \tau_1(y_0) - 2d$. A parameter vector is recovered from its actual valuations. Across different seeds, distinct full first-hitting depths are a sufficient distinctness certificate. $\square$

Theorem C.1 (exact count recurrence).

Set

$$N_r(T) = \#\{k \in \mathbb{Z}_{\geq 0}^r : \sum_{i=1}^r 3^{i+1} k_i \leq T\},$$

with $N_r(T) = 0$ for $T < 0$ and $N_0(T) = 1$ for $T \geq 0$. Then

$$\boxed{N_r(T) = N_{r-1}(T) + N_r(T - 3^{r+1})}. \quad (\text{C.rectcount})$$

The number with last nonzero coordinate exactly r is

$$\Delta_r(T) = N_r(T - 3^{r+1}). \quad (\text{C.rectlayer})$$

Proof. Partition by $k_r = 0$ or $k_r > 0$. In the second part subtract one from k_r ; this is a bijection to the indicated lower-budget set. $\square$

For $T \geq 0$, the unit-cube bounds give

$$\frac{T^r}{r! 3^{r(r+3)/2}} \leq N_r(T) \leq \frac{(T + \sum_{i=1}^r 3^{i+1})^r}{r! 3^{r(r+3)/2}}. \quad (\text{C.volume})$$

Flooring points of the real simplex proves the lower bound via unit cubes; its cube cover lies inside the enlarged simplex for the upper bound. The resulting asymptotic is for fixed r , not uniformly growing dimension.

Theorem C.2 (retention of sharper earlier layers).

Unrolling the original prefix gives

$$L_0^{(r)} = y_0 + \sum_{i=1}^r \frac{2^{A_{i-1}}}{3^i}, \quad \beta_{r+1} - \beta_r = \frac{2^{e+A_r}}{4^d 3^{r+1}} > 0. \quad (\text{C.betastep})$$

Thus replacing each retained block ceiling by the final β_r can weaken a coarse certificate. Instead define

$$G_R(T) = \alpha_0^{-1} + \sum_{r=1}^R \frac{N_r(T - 3^{r+1})}{\beta_r}. \quad (\text{C.rectscore})$$

Then

$$\boxed{G_R(T) - G_{R-1}(T) = N_R(T - 3^{R+1})/\beta_R.} \quad (\text{C.rectincrement})$$

The right side is positive exactly when that new rectangular layer is nonempty.

Proof. The zero vector gives the base target block and the factor α_0^{-1} . A nonzero vector has a unique last nonzero coordinate r ; its later coordinates vanish, so the target is exactly its r -prefix construction, with coefficient less than β_r . Apply (R.blockcomparison) separately to those Δ_r blocks and retain their budgets. $\square$

Equation (C.rectincrement) gives the exact increment of the guaranteed multiplier.

Exact evaluation without canonical root enumeration The same counts admit the following exact radix evaluation. Let

$$H_{r,p}(n) = [z^n](1 - z)^{-(p+1)} \prod_{j=0}^{r-1} (1 - z^{3^j})^{-1}.$$

Then $N_r(T) = H_{r,0}(\lfloor T/9 \rfloor)$, $H_{0,p}(n) = \binom{n+p}{p}$, and all values with $n < 0$ vanish. For $n = 3m + a$, $a \in \{0, 1, 2\}$, define

$$c_j(p, a) = \sum_{v=0}^j (-1)^v \binom{j}{v} \binom{-3(v+1) + a + p + 1}{p+1}.$$

Using generalized binomial coefficients for negative upper arguments,

$$H_{r,p}(3m + a) = \sum_{j=0}^{p+1} c_j(p, a) H_{r-1,j}(m). \quad (\text{C.radix})$$

To prove it, expand the polynomial $\binom{3u+a+p+1}{p+1}$ in the basis $\binom{u+j}{j}$, whose coefficients are the displayed backward differences. Apply the generating-function identity for division of the geometric weights by three. This is a signed-integer recurrence; no decimal values decide the counts.

The canonical rectangular prefixes (7, 15, 14) give

$$E_b = 1 + G_{101,7}(2b - 4) + G_{193,15}(2b - 4) + G_{12289,14}(2b - 7). \quad (\text{C.E})$$

It remains a separately named benchmark in the canonical comparison below.

C.3. General increments, heights, residues and clocks

Now allow **arbitrary nonnegative integer increments** $\ell_1, \dots, \ell_r$, not necessarily even. Write

$$S = \sum_{i=1}^r \ell_i, \quad S_j = \sum_{i=1}^j \ell_i.$$

Define, whenever the divisions are integral,

$$Y_r = y_r, \quad Y_{i-1} = \frac{2^{a_i + \ell_i} Y_i - 1}{3}, \quad (\text{C.redistribution})$$

and let the incoming root exponent be

$$a_{\text{in}}(s) = 2(s - d) + e - S, \quad M_s = \frac{2^{a_{\text{in}}(s)} Y_0 - 1}{3}. \quad (\text{C.incoming})$$

Here e is the original seed parameter. It need not equal $e(Y_0)$. Correct incoming parity will follow from the compensated congruence; one must not substitute a nonintegral half-shift into the one-step constructor formula.

Assume the backward values are positive odd units and satisfy

$$2^e Y_0 \equiv 4^d 2^S \pmod{9}. \quad (\text{C.compensated})$$

The residue-completion theorem gives a complete deterministic way to meet these conditions.

Lemma C.3 (geometry and exact affine coefficient).

For $0 \leq j < r$,

$$2^{S-S_j} y_j \leq Y_j < L_j^{(r)} 2^{S-S_j}. \quad (\text{C.geometry})$$

Consequently

$$\alpha(\ell) = \frac{2^e Y_0}{4^d 2^S}, \quad \alpha_0 \leq \alpha(\ell) < \beta_r, \quad M_s = \alpha(\ell) R_s + \frac{\alpha(\ell) - 1}{3}. \quad (\text{C.affine})$$

The exact subtraction terms are

$$\alpha(\ell) = \beta_r - \frac{2^e}{4^d} \sum_{i=1}^r \frac{2^{A_{i-1}}}{3^i 2^{\ell_i + \dots + \ell_r}}. \quad (\text{C.exactalpha})$$

Proof. Backwards induction proves the lower bound. If $Y_i \geq 2^{S-S_i} y_i$, then

$$Y_{i-1} \geq 2^{S-S_{i-1}} y_{i-1} + \frac{2^{S-S_{i-1}} - 1}{3} \geq 2^{S-S_{i-1}} y_{i-1}.$$

Dropping the subtractive one in each inverse step proves the strict upper bound; already the last inverse step is strictly below its leading term. Unrolling the complete affine composition yields (C.exactalpha). Dividing (C.geometry) at $j = 0$ by the stated powers gives (C.affine). $\square$

Lemma C.4 (height correction).

Assume $\beta_r < 1$. There is no need to impose $\Gamma_r < 1$. Define the least nonnegative integer

$$\chi_r = \min\{v \geq 0 : \Gamma_r 2^{e-2} \leq 4^v\}, \quad T_r = 2b - d + e - 1 - \chi_r. \quad (\text{C.heightdebit})$$

Restrict the total increment by

$$S \leq 2T_r. \quad (\text{C.budget})$$

If $T_r < 0$, use no nonzero vectors at this layer. Under the base conditions (C.seed)–(C.baseheight), every selected root is at least 16^b , every subsequent modified odd state is strictly below 16^b , and the incoming exponent is positive.

Proof. At $s \geq H$,

$$a_{\text{in}}(s) \geq 2H - 2d + e - 2T_r = 4 - e + 2\chi_r \geq 2.$$

The lower bound $\alpha \geq \alpha_0$ retains the finite minimum-root height condition. On the other hand,

$$Y_j < L_j^{(r)} 2^{2T_r} \leq \Gamma_r 2^{e-2} 4^{-\chi_r} 16^b \leq 16^b. \quad (\text{C.height})$$

The inequality is strict because the leading-term bound is strict. The unchanged tail already lies below 16^b . $\square$

Thus a large intermediate leading coefficient costs a fixed number of budget units rather than automatically prohibiting the prefix. Earlier layers keep their own χ_j, T_j, β_j ; a later height correction must not be imposed on them.

Lemma C.5 (residues, distinctness, and clock).

Under (C.compensated), the coefficient α is one modulo nine in the dyadic-rational sense. Therefore $r \mapsto \alpha r + (\alpha - 1)/3$ is an invertible affine map modulo every 3^q , with constant term divisible by three. The baseline indices $\mathcal{S}_q$ give one root in every unit residue, unchanged.

For every selected vector,

$$\tau_1(Y_0) = \tau_1(y_0) + S, \quad \boxed{\tau_1(M_s) = 2s + e + 1 + \tau_1(y_0) - 2d.} \quad (\text{C.clock})$$

Proof. Multiplying (C.compensated) by the inverse of $4^d 2^S$ proves the coefficient congruence. The root-period lemma for R_s , followed by the affine permutation, gives complete coverage.

Oddness of each Y_i makes its valuation exactly $a_i + \ell_i$. The lower geometry bound implies $Y_j \geq y_j > 1$ for $j < r$, so there is no premature hit of one. The unchanged tail follows from $Y_r = y_r$. Sum the ordinary step costs and cancel S against the incoming exponent in (C.incoming). $\square$

Within one seed, the full vector of increments, padded by zeros to the original first-hitting depth, is recovered from the actual valuation word. Different vectors give different targets. To combine different seeds, require distinct perturbed targets; distinct original full first-hitting depths are a sufficient certificate, used in the canonical example. The height barrier then proves the joint root antichain, including the unchanged baseline block.

C.4. Residue completion

The necessary congruences can be solved backwards from the certified endpoint, with a bounded correction at each step.

Theorem C.6 (complete residue-completion algorithm).

Fix a prefix of length $r \geq 1$, with base conditions (C.seed), and choose

$$(k_1, \dots, k_r) \in \mathbb{Z}_{\geq 0}^r, \quad (\varepsilon_1, \dots, \varepsilon_{r-1}) \in \{1, 2\}^{r-1}.$$

The labels ε_i prescribe the unit residues of the intermediate states Y_i . They are deterministic branch labels, not probabilities.

Start with $Y_r = y_r$. For $i = r, r-1, \dots, 2$, choose the unique $c_i \in \{0, \dots, 5\}$ such that

$$2^{a_i+c_i} Y_i \equiv 3\varepsilon_{i-1} + 1 \pmod{9}, \quad \ell_i = 6k_i + c_i, \quad Y_{i-1} = \frac{2^{a_i+\ell_i} Y_i - 1}{3}. \quad (\text{C.internal})$$

Put $S' = \sum_{i=2}^r \ell_i$, with $S' = 0$ when $r = 1$, and form

$$D \equiv 2^{e+a_1} Y_1 - 3 \cdot 4^d 2^{S'} \pmod{27}. \quad (\text{C.topresidue})$$

Choose the unique $c_1 \in \{0, \dots, 17\}$ with

$$2^{c_1} D \equiv 2^e \pmod{27}, \quad \ell_1 = 18k_1 + c_1, \quad Y_0 = \frac{2^{a_1+\ell_1} Y_1 - 1}{3}. \quad (\text{C.top})$$

This always constructs positive odd unit states, exact valuations $a_i + \ell_i$, and the compensated congruence (C.compensated).

Moreover, the construction is a bijection onto the nonnegative increment vectors for which all the backward states are positive odd units and (C.compensated) holds. No coordinatewise condition $3^{i+1} \mid \ell_i/2$ is required.

Proof: internal steps. Powers of two have order six modulo nine and enumerate all its units. Since Y_i is a unit, exactly one exponent residue gives each of the two target numerator residues 4 and 7. Division by

three gives $Y_{i-1} \equiv \varepsilon_{i-1} \pmod{3}$. The quotient is integral and odd. Positivity and the lower bound by the original state follow from the induction in Lemma C.3.

Proof: first step and root condition. The residue D is a unit modulo three and hence modulo 27. Powers of two have order eighteen modulo 27 and enumerate all units: $2^9 \equiv -1 \pmod{27}$, while $2^6 \not\equiv 1 \pmod{27}$ and $2^2 \not\equiv 1 \pmod{27}$, excluding the proper candidate orders after using the nine-step sign. Thus (C.top) has exactly one residue solution.

Reducing (C.top) modulo three gives $2^{a_1+\ell_1}Y_1 \equiv 1 \pmod{3}$, so Y_0 is integral. In full precision it reads

$$2^e(2^{a_1+\ell_1}Y_1 - 1) \equiv 3 \cdot 4^d 2^{S'+\ell_1} \pmod{27}.$$

Divide this congruence by three, using the integer numerator, to obtain

$$2^e Y_0 \equiv 4^d 2^S \pmod{9}.$$

It also implies $3 \nmid Y_0$. Positive oddness follows as before.

Proof: bijection. Conversely, take any vector satisfying the stated physical and congruence conditions. For $i \geq 2$, the residue of Y_{i-1} determines ε_{i-1} , and $\ell_i \pmod{6}$ must be the unique solution of (C.internal). Its quotient is $k_i = \lfloor \ell_i/6 \rfloor$. The final compensated congruence is equivalent to (C.top), so $\ell_1 \pmod{18}$ is the unique first residue and $k_1 = \lfloor \ell_1/18 \rfloor$. These recovered inputs reproduce the vector. $\square$

For $r = 1$, one has $D = 2^e + 3(2^e y_0 - 4^d) \equiv 2^e \pmod{27}$, so the base congruence makes $c_1 = 0$. Thus $\ell_1 = 18k_1$, exactly the original single-parameter lift. For a general prefix,

$$S = 18k_1 + 6 \sum_{i=2}^r k_i + \sum_{i=1}^r c_i. \quad (\text{C.totalincrement})$$

The uniformly bounded residue corrections give the counting cost.

A small parity-changing example Use the seed 101, with $d = 4, e = 1$, and its first two valuations $(4, 1)$, ending at 29. Set $k = (0, 0)$ and choose $Y_1 \equiv 2 \pmod{3}$. The algorithm gives

$$(\ell_1, \ell_2) = (11, 2), \quad Y_1 = 77, \quad Y_0 = 841045.$$

Indeed

$$3 \cdot 841045 + 1 = 2^{15} \cdot 77, \quad 3 \cdot 77 + 1 = 2^3 \cdot 29.$$

The total added valuation is 13. The roots are

$$M_s = (841045 \cdot 2^{2s-20} - 1)/3, \quad \tau_1(M_s) = 2s + 19.$$

For example, at $b = 5, H = 11, s = 11$, the root is $1121393 > 16^5$, while all its subsequent odd states are below 16^5 . Its ordinary first hitting time is 41. This is an arithmetic example only; the density theorem still retains its canonical analytic range for b .

Here $e(Y_0) = 2$, although the seed parameter is $e = 1$. The compensated incoming exponent has automatically changed parity. Keeping $e(Y_0) = e(y_0)$ would wrongly exclude this valid family.

Relation to the rectangular family Every rectangular vector satisfies the completed family's conditions: it is integral, all intermediate states are units, and its affine coefficient is one modulo nine. Theorem C.6 therefore recovers it uniquely. The completed family is strictly larger; the example above contains an odd increment and cannot be a rectangular vector.

The original recurrence (C.rectcount) remains the exact count of its original rectangular set. It must not be presented as the exact count of the larger residue-completed set.

C.5. Counting completed layers

Lemma C.7 (quotient-simplex count).

For $r \geq 1$, define

$$Q_r(n) = \#\{k \in \mathbb{Z}_{\geq 0}^r : 3k_1 + k_2 + \cdots + k_r \leq n\}, \quad Q_r(n) = 0 \ (n < 0).$$

Then

$$Q_r(n) = \sum_{v=0}^{\lfloor n/3 \rfloor} \binom{n-3v+r-1}{r-1}. \quad (\text{C.quotientcount})$$

This counts all choices of the remaining $r-1$ coordinates and unused budget after choosing $k_1 = v$.

For an efficient exact evaluation, write $n = 3m + a$, $a \in \{0, 1, 2\}$, and

$$P_{r,a}(u) = \binom{3u+a+r-1}{r-1}, \quad d_j = \sum_{v=0}^j (-1)^{j-v} \binom{j}{v} P_{r,a}(v).$$

Then

$$Q_r(3m+a) = \sum_{j=0}^{r-1} d_j \binom{m+1}{j+1}. \quad (\text{C.quotienteval})$$

Proof. Reindex the sum in (C.quotientcount) as $\sum_{u=0}^m P_{r,a}(u)$. The polynomial has degree $r-1$, so its finite Newton expansion is $P(u) = \sum_{j=0}^{r-1} d_j \binom{u}{j}$. Summing and using $\sum_{u=0}^m \binom{u}{j} = \binom{m+1}{j+1}$ proves (C.quotienteval). All arithmetic is integer arithmetic. $\square$

Thus the dimension-43 count is evaluated from a degree-42 polynomial using exact binomial operations.

Lemma C.8 (disjoint last-nonzero layers with explicit correction budgets).

A nonzero increment vector has a unique last nonzero coordinate r . Later coordinates are zero, so it is the r -prefix construction. At its last internal step, $Y_r = y_r$. The two residue completions for c_r are

$$c_r = 0 \quad \text{or} \quad c_r = \sigma_r := \begin{cases} 2, & y_{r-1} \equiv 1 \pmod{3}, \\ 4, & y_{r-1} \equiv 2 \pmod{3}. \end{cases}$$

The first choice preserves the original residue of y_{r-1} , and the other switches it.

For $r \geq 2$, each of the earlier $r-2$ internal residues is freely chosen from two units. The first correction is at most 17, and those earlier internal corrections are at most five each. Consequently a fixed last correction $c \in \{0, \sigma_r\}$ has total correction at most

$$\delta_{r,c} = 17 + 5(r-2) + c = 5r + 7 + c. \quad (\text{C.correctionbudget})$$

To ensure $\ell_r > 0$, use three disjoint classes:

Class	Last residue correction	Last quotient	Earlier branch choices	Guaranteed free-quotient budget
A	0	$k_r \geq 1$	2^{r-2}	subtract $\delta_{r,0} + 6$ from $2T$
B	σ_r	$k_r \geq 1$	2^{r-2}	subtract $\delta_{r,\sigma_r} + 6$ from $2T$
C	σ_r	$k_r = 0$	2^{r-2}	subtract δ_{r,σ_r} from $2T$

Subtract one from k_r in A and B. In C omit the last quotient. The exact number in the selected parameter classes is

$$\Lambda_r(T) = 2^{r-2} \left[Q_r \left(\left\lfloor \frac{2T - \delta_{r,0} - 6}{6} \right\rfloor \right) + Q_r \left(\left\lfloor \frac{2T - \delta_{r,\sigma_r} - 6}{6} \right\rfloor \right) + Q_{r-1} \left(\left\lfloor \frac{2T - \delta_{r,\sigma_r}}{6} \right\rfloor \right) \right] \quad (r \geq 2), \quad (\text{C.selectedcount})$$

where $\Lambda_r(T) = 0$ for $T < 0$, and

$$\Lambda_1(T) = \max(0, \lfloor T/9 \rfloor).$$

Proof. The last-step residues follow by multiplying the original numerator residue, either four or seven modulo nine, by powers of two. The alternate valid unit output occurs at exponent two or four respectively. At this last step, c_r is independent of its own quotient k_r , because $2^6 \equiv 1 \pmod{9}$. Earlier corrections may depend on downstream quotients through the computed states; the uniform bounds cover that dependence.

For each fixed earlier branch string, (C.totalincrement) and the correction budget ensure $S \leq 2T$. The quotient counts are precisely (C.quotientcount). Theorem C.6 is injective in all the choices, and the three last-quotient/residue classes are disjoint. Therefore multiplication by 2^{r-2} and addition are exact counts of the chosen parameter classes. At $r = 1$, the nonzero shifts are exactly $18k_1 \leq 2T$, $k_1 \geq 1$. $\square$

These are exact counts of **explicitly selected safe parameter sets**. They are lower bounds on the count of all admissible completed vectors, because an individual vector can have a smaller correction than the common budget used above. Neither rounding nor a distributional assumption enters the selected counts.

The larger finite domain For every prefix r , let $\mathcal{L}_r$ consist of all nonnegative vectors with last nonzero coordinate r , satisfying the physical unit conditions, (C.compensated), and $S \leq 2T_r$. This is a finite, deterministically decidable set. Theorem C.6 enumerates it by bounded quotient vectors and unit branches, followed by its literal total-budget test. It contains the selected classes of Lemma C.8.

Padded zero suffixes identify these as disjoint layers of the full original certificate. The single all-zero vector is added once per seed, not once for every prefix.

For a depth- D seed, all these choices inject into the unrestricted depth- D completion parameters with $18k_1 + 6 \sum_{i \geq 2} k_i \leq 2T_0$. Thus the number of its blocks, including its zero vector, is at most

$$2^{D-1} Q_D(\lfloor T_0/3 \rfloor). \quad (\text{C.blockcount})$$

This upper bound deliberately ignores some completion and height costs. It is used only for the domain condition.

Structural count improvement at fixed dimension For fixed r , let $\mathcal{N}_r(T)$ count all completed vectors with $S \leq 2T$, before imposing a prefix-specific height correction. For $r \geq 2$, the last residue correction is at most four, the first is at most seventeen, and the other corrections total at most $5(r-2)$. Thus

$$2^{r-1} Q_r \left(\left\lfloor \frac{2T - (5r + 11)}{6} \right\rfloor \right) \leq \mathcal{N}_r(T) \leq 2^{r-1} Q_r(\lfloor T/3 \rfloor).$$

The same quotient-simplex argument as for (C.volume), now with weights $(3, 1, \dots, 1)$, gives $Q_r(n) = n^r / (3r!) + O_r(n^{r-1})$. Therefore

$$\mathcal{N}_r(T) = \frac{2^{r-1}}{3^{r+1}r!} T^r + O_r(T^{r-1}).$$

For $r = 1$ and $T \geq 0$, the exact count is $1 + \lfloor T/9 \rfloor$, with the same leading coefficient. The selected last-layer formula (C.selectedcount) has this same degree- r leading coefficient: the first two terms supply it, while the third and the shorter layers have lower degree.

Relative to the rectangular lattice's degree- r coefficient in (C.volume), the ratio is

$$2^{r-1} 3^{(r-1)(r+2)/2}.$$

This asymptotic comparison holds with r fixed. The canonical finite result uses the explicit counts and inequalities below.

Retained-layer multiplier for the completed construction Define

$$\widehat{G}_R = \alpha_0^{-1} + \sum_{r=1}^R \frac{\Lambda_r(T_r)}{\beta_r}. \quad (\text{C.layerscore})$$

Then

$$\widehat{G}_R - \widehat{G}_{R-1} = \Lambda_R(T_R) / \beta_R \geq 0.$$

Every nonempty selected new layer gives a strictly positive increment. Crucially, a later β_R , height correction, or shorter budget never replaces the earlier layers' retained certificates. This is the analogue of (C.rectincrement), with both residue corrections and height bounds included.

Using the exact count of $\mathcal{L}_r$, or finer per-vector coefficients from (C.exactalpha), can improve this convenient explicit multiplier further. The selected-set formula is not asserted to equal the exact shared score.

C.6. Finite-family score comparison

Theorem C.9 (residue-completed finite-family score comparison).

Retain the analytic and counting inputs of the main text. Fix a nonempty finite collection of certified convergent odd unit seeds y , greater than one, each with parameters $e = e(y), d$ satisfying (C.seed)–(C.baseheight). For each seed choose a terminal prefix length $1 \leq R \leq D$, where D is its certified first-hitting depth, and assume $\beta_R < 1$. Earlier leading coefficients are smaller by (C.betastep). Form χ_r, T_r separately for every $1 \leq r \leq R$.

Use the zero vector once and any or all of the finite completed layers $\mathcal{L}_r$. In particular one may use the full layer sets. Require that generated targets from different seeds be distinct; pairwise distinct full accelerated first-hitting depths suffice. Retain the original block $\{R_s : s \in \mathcal{S}_q\}$.

The finite score comparisons below hold for every $q \geq 2$. For the persistent-mark and density conclusions, take $q = q_N$ and $p = P$, as in (R.markapplication). The resulting root family is a finite convergent antichain, partitioned into complete unit-residue blocks; at $q = q_N$ these blocks carry the common persistent marks (G.marks). With the selected classes included, set

$$\mathfrak{A} = 1 + \sum_{\text{seeds } y} \widehat{G}_{R_y}.$$

Then

$$\boxed{W_{\mathcal{F}}^{\text{joint}}(q, P) > \mathfrak{A} W(q, P) > 0.} \quad (\text{C.master})$$

All roots are at least 16^b , every later odd state from a new root is below 16^b , the largest root is the largest baseline root, and every root clock is given by (C.clock).

Let

$$J_{\text{up}} = 1 + \sum_y 2^{R_y-1} Q_{R_y}(\lfloor T_{0,y}/3 \rfloor).$$

The following bound holds:

$$W_{\mathcal{F}}^{\text{joint}} < \frac{4}{3} J_{\text{up}} 16^{-b}. \quad (\text{C.domain})$$

At $q = q_N$, whenever the shared score is within W_{max} , (D.finitepool) gives the corresponding every-clock density bound. If the right side of (C.domain) is at most 2^{-4096} , all positive scores in (C.master) are also in the cubic comparison range.

Proof. The geometric and residue lemmas prove each block's admissibility. A target's modified first-hitting depth is unchanged, so the stated cross-seed condition proves target distinctness. Within a seed, actual valuations recover the increment vector and hence its unique layer. A root's next odd value identifies its block. The height barrier proves no root subsequently visits any selected root. Therefore there is one aggregate $1/x$ bound on each source weight.

Every block uses the same capacity B_q and mass requirement $K_q P$. At $q = q_N$, the mark identity, uniform over roots, and full variation bound supply these as actual persistent marks in every complete block. For the zero vector use the factor α_0^{-1} . In layer r , (C.affine) gives $M_s < \beta_r R_s$, so (R.blockcomparison) supplies a strict block factor β_r^{-1} . Count the disjoint selected classes by (C.selectedcount), sum their budgets, retain the baseline block, and apply the shared-score dominance (R.jointdominance). This proves (C.master).

For (C.domain), the constant vector $v_r = P \leq 1$ is feasible in (R.jointscore). Each block obeys $M_{s+1} = 4M_s + 1$ on the full sequence of consecutive indices and starts at least at 16^b . Its finite unit-subset reciprocal sum is strictly less than $\sum_{j \geq 0} 4^{-j} 16^{-b} = (4/3) 16^{-b}$. Bound the number of blocks by (C.blockcount) and include the baseline block. This proves an upper bound on the finite allocation minimum, with no sum over infinitely many physical histories. The remaining density conclusion is (D.finitepool). $\square$

For a seed used only through a proper prefix, (C.blockcount) and (C.domain) use that prefix length and pad shorter layers only to it. Distinctness still uses the full first-hitting depth of its certified original trajectory. The canonical application below uses full prefixes, so these lengths coincide there.

Strongest finite certificate retained by the theorem For the explicitly defined finite family, (R.jointscore) retains the full residue coupling and exact root denominators. It dominates the sum of individual block minima, which in turn dominates the exact coefficient sum (R.blockcomparison) and the selected-count multiplier (C.master). The resulting certificate hierarchy is

$$\begin{aligned} \text{shared residue allocation} &\geq \text{separate block allocations} \\ &> \text{exact affine-coefficient comparison} \\ &\geq \text{certified cell/layer comparisons.} \end{aligned}$$

The rectangular E_b certificate is retained as a benchmark. In the canonical completed family its rectangular blocks are contained in the full completed layers: their original prefixes have zero height correction, their budgets are unchanged, and the completion theorem recovers their vectors. The larger family is not obtained by changing the definition of E_b .

C.7. Canonical certificates and density bounds

Here $q = q_N$, $p = P$, and $W_{\mathcal{F}}^{\text{joint}} = W_{\mathcal{F}}^{\text{joint}}(q_N, P)$. The three certificates used here have the following exact data:

Seed y	d	e	Full odd depth D	Ordinary time	Maximum odd state	α_0
101	4	1	7	25	101	101/128
193	5	2	43	119	3077	193/256
12289	8	2	14	50	12289	12289/16384

Their complete valuation words, in forward order and ending at the first hit of one, are:

101:

4, 1, 3, 1, 2, 3, 4

193:

2, 2, 3, 2, 1, 1, 1, 1, 2, 2, 1, 2, 1, 1, 2,
1, 1, 1, 2, 3, 1, 1, 2, 1, 2, 1, 1, 1, 1, 1,
3, 1, 1, 1, 4, 2, 2, 4, 3, 1, 1, 5, 4

12289:

2, 2, 2, 2, 2, 4, 1, 5, 3, 3, 1, 2, 3, 4

Apply $3y + 1$, check the displayed exact valuation, and divide by its power of two at every step. This reconstructs each finite certificate; no assumption about an uncertified trajectory is needed.

The full-prefix leading data are

Seed	A_D	β_D	Γ_D	χ_D	Final T_D at $b = 2^{25}$
101	18	2048/2187	2048/2187	0	67,108,860
193	76	$2^{68}/3^{43}$	262144/19683	2	67,108,858
12289	36	4194304/4782969	4194304/4782969	0	67,108,857

All full-prefix β_D are less than one, and hence so are the earlier β_r . For 193, Γ_{43} lies strictly between four and sixteen, so the least height correction is exactly two. Its earlier layers retain the smaller correction appropriate to each prefix.

The base minimum-root excesses over 16^b are, respectively,

$$\frac{5 \cdot 16^b - 32}{96}, \quad \frac{16^b - 64}{192}, \quad \frac{16^b - 4096}{12288}.$$

They are positive for $b \geq 4$, and the maximum certified odd state of each target is below 16^4 . These verify the finite base conditions. The analytic application continues to use $b = 2^{25}$.

Exact score multiplier For every seed and prefix, compute $\beta_{y,r}, \Gamma_{y,r}, \chi_{y,r}$ from (C.leading) and (C.heightdebit), and $\Lambda_{y,r}$ from (C.selectedcount). Define

$$\mathfrak{A}_b = 1 + \sum_{(y,d,D)=(101,4,7),(193,5,43),(12289,8,14)} \left[\frac{1}{\alpha_y} + \sum_{r=1}^D \frac{\Lambda_{y,r}(2b-d+e(y)-1-\chi_{y,r})}{\beta_{y,r}} \right]. \quad (\text{C.factor})$$

This is an exact rational definition involving 64 prefix layers. Each integer count is evaluated by the finite-difference formula (C.quotienteval). It does not require a root list, a residue table at the canonical q_N , or an optimized value of c_{FM} .

At $b = 2^{25}$,

$$\begin{aligned} \mathfrak{A}_b &\approx 2.9273409933096237 \times 10^{275}, \\ \mathfrak{A}_b^3 &\approx 2.5085337209674080 \times 10^{826}, \\ (\mathfrak{A}_b/E_b)^3 &\approx 2.0173520932097131 \times 10^{697}. \end{aligned}$$

The following strict inequalities

$$\mathfrak{A}_b^3 > \frac{250}{100} 10^{826}, \quad (\mathfrak{A}_b/E_b)^3 > \frac{201}{100} 10^{697} \quad (\text{C.numericguards})$$

are checked by integer cross-multiplication of exact fractions. Decimal displays play no role in the inequalities.

Benchmark root families

Certificate	Fixed meaning	Guaranteed score multiplier over W_*
C_b	One-coordinate period-nine lifts	$57101167060327/1953295 \approx 2.9233253 \times 10^7$
D_b	Rectangular multivariate volume estimate, prefixes (7, 13, 13)	$5.324667695451 \dots \times 10^{42}$
E_b	Exact rectangular counts with retained layers, prefixes (7, 15, 14)	$1.0753406914595744 \dots \times 10^{43}$
$\mathfrak{A}_b$	Residue-completed increments, height corrections, full prefixes (7, 43, 14)	$2.9273409933096237 \dots \times 10^{275}$

For clarity, the one-coordinate and rectangular-volume benchmarks are the literal quantities

$$\begin{aligned} C_b &= 1 + \left(1 + \left\lfloor \frac{2b-4}{9} \right\rfloor \right) \left(\frac{24}{19} + \frac{192}{145} \right) + \left(1 + \left\lfloor \frac{2b-7}{9} \right\rfloor \right) \frac{12288}{9217}, \\ D_b &= 1 + \frac{(2b-4)^7}{7! 2^{11} 3^{28}} + \frac{(2b-4)^{13}}{13! 1269760 3^{91}} + \frac{(2b-7)^{13}}{13! 1310720 3^{91}}. \end{aligned} \quad (\text{C.earlier})$$

The affine comparison and the rectangular simplex bound give these weaker benchmarks by retaining their indicated blocks and leading ceilings. The definition of E_b in (C.E) retains each earlier layer's sharper coefficient. The full completed family contains those rectangular blocks: their original prefixes have zero height correction and unchanged budgets. None of these symbols is redefined to mean the larger completed family. The exact counts and rational evaluations accompany the computational evidence.

Domain and root clocks The canonical block-count upper bound is

$$J_{\text{up}} = 1 + 2^6 Q_7(\lfloor 67108860/3 \rfloor) + 2^{42} Q_{43}(\lfloor 67108860/3 \rfloor) + 2^{13} Q_{14}(\lfloor 67108857/3 \rfloor).$$

The exact scalar check is

$$4J_{\text{up}} < 3 \cdot 2^{916}. \quad (\text{C.countguard})$$

Equations (C.master)–(C.domain) therefore give

$$0 < \mathfrak{A}_b W_* < W_{\mathcal{F}}^{\text{joint}} < 2^{916-4b} = 2^{-134216812} < 2^{-4096}. \quad (\text{C.domainguard})$$

The feasible constant vector supplies this upper bound for the allocation minimum.

The large relative factor has a separate absolute scale. All retained mass bounds are at most W/μ , including (D.refinedmass), so

$$c_{\text{comp}}(W) \leq \frac{32}{225}(e^{9W/4} - 1) \leq \frac{16}{25}W < W \quad (0 < W \leq 2^{-4096}),$$

using $e^z - 1 \leq z/(1-z) \leq 2z$ for $0 \leq z \leq 1/2$. Thus the strengthened **certified formula itself** remains smaller than $2^{-134216812}$. The displayed power is a small-score domain condition, not an evaluation of the absolute density formula. The further level loss is quantified for the named formulas in (Q.formulascale) and (Q.refinedscale). Neither statement bounds the actual convergence density from above.

The ordinary root-clock excess over the indexed baseline root R_s , whose clock is $2s + 1$, is

$$e(y) + \tau_1(y) - 2d = \begin{cases} 18, & y = 101, \\ 111, & y = 193, \\ 36, & y = 12289. \end{cases}$$

The baseline block is retained, and every seed's zero block is retained. Thus the largest root is exactly the largest baseline root, and the largest root clock is exactly the old maximum plus 111. None of the quotient vectors, residue choices, prefix lengths, or height corrections changes these formulas.

Consequence for the density formula Apply (D.finitepool) to the completed family with target one, and apply (Q.cubicscore) directly to the pairs $(W_*, \mathfrak{A}_b W_*)$, $(E_b W_*, \mathfrak{A}_b W_*)$, and $(\mathfrak{A}_b W_*, W_{\mathcal{F}}^{\text{joint}})$. Their hypotheses follow from (C.domainguard), (C.numericguards), and positivity of E_b . With the constants named in (I.constants), this proves (I.rootdensity). For every $c > c_0$,

$$\begin{aligned} \underline{d}(\mathcal{G}_c) &\geq c_{\text{FM}}(W_{\mathcal{F}}^{\text{joint}}) > c_{\text{FM}}(\mathfrak{A}_b W_*) > 2.50 \times 10^{826} c_{\text{FM}}(W_*), \\ c_{\text{FM}}(\mathfrak{A}_b W_*) &> 2.01 \times 10^{697} c_{\text{FM}}(E_b W_*). \end{aligned} \quad (\text{C.densitygain})$$

These compare certified lower-bound formulas at the same analytic envelope and moment tables. They do not compare actual densities or evaluate the absolute optimized constant. The second line follows from the literal score pair, not division of independent lower bounds. The primary quantitative improvement of this appendix is the score guarantee (C.master).

Within these three fixed prefixes and incoming budgets, replacing every selected count and every reciprocal coefficient by its allowed upper bound changes the additive comparison factor by at most

$$\frac{1 + (4/3)(J_{\text{up}} - 1)}{\mathfrak{A}_b} = 1.1989376587603659 \dots < 1.199. \quad (\text{C.localceiling})$$

Every block coefficient exceeds $3/4$, so each contributes less than $4/3$, and (C.blockcount) bounds the number. This bounds that additive coefficient comparison only. It does not bound the shared allocation minimum, other seeds, or a changed height geometry.

The combined refined bound is $\underline{d}(\mathcal{G}_c) \geq c_{\text{comp}}(W_{\mathcal{F}}^{\text{joint}})$ for every $c > c_0$, by (D.finitepool). Equation (Q.sixteengain) makes it strictly larger than the pure-fractional bound displayed here. The cubic multipliers in (C.densitygain) retain their named c_{FM} meaning; (Q.sixteencubic) supplies the corresponding comparison for $c_{\text{FM},16}$, not for c_{comp} .

C.8. A shared-allocation example: the 101 block

There is a strict gain beyond summing separate minima even before the valuation extension. For the base block

$$M_s = \frac{101}{128}R_s - \frac{9}{128}, \quad \phi(r) = 128^{-1}(101r - 9) \pmod{3^q},$$

the affine permutation sends the baseline root residue to the new one. Writing $\alpha = 101/128$, the valuation identity gives

$$\phi^k(r) - r = (\alpha^k - 1)(r + 1/3), \quad \nu_3(\phi^k(r) - r) = 2 + \nu_3(k) \quad (k \geq 1). \quad (\text{C.101orbits})$$

Indeed $\nu_3(\alpha - 1) = 3$; cubing a unit congruent to one modulo three raises the valuation of its difference from one by one, and an exponent prime to three preserves it. Also $\nu_3(r + 1/3) = -1$ for every integer r . Thus, for $q \geq 3$, every unit orbit has length 3^{q-2} .

For $q \geq 4$, $0 < p \leq 1$, both separate allocation minima uniquely fill the largest roots first, to the common cap $B_q = (2/3)2^q$, with at most one partial entry. This follows by transferring mass from a larger reciprocal cost to a smaller one; all costs are distinct. Their minimizing index vectors therefore agree. A common residue-labeled minimizer would have to be ϕ -invariant. Its nonempty support has size

$$\left\lceil \frac{K_q p}{B_q} \right\rceil \leq \lceil (3/2)^q \rceil < 3^{q-2}.$$

For the last inequality, $(3/2)^q/3^{q-2} = 9/2^q \leq 9/16$ and $3^{q-2} \geq 9$. Such a support cannot contain a whole orbit. By (R.jointdominance),

$$W_{\{R,M\}}^{\text{joint}}(q,p) > W(q,p) + W_{101}(q,p) > (1 + 128/101)W(q,p). \quad (\text{C.strict101})$$

No invariance of g_N under ϕ is used. The strict conclusion excludes $p = 0$; at $q = 1, 2$ the affine map is the identity. This argument supplies strictness without a numerical lower bound on the extra shared-allocation gap.

C.9. Limits at fixed height

The completed construction remains in a specific height geometry. Put $X = 16^b$, with the canonical $b = 2^{25}$. Its affine blocks have

$$M_s = 4^{s-H}n + \frac{4^{s-H} - 1}{3}, \quad X \leq n \leq R_H = \frac{4X - 1}{3}.$$

To preserve exactly the baseline unit index set, one requires $n \equiv 2 \pmod{3}$, since $H \equiv 2 \pmod{3}$. For an odd n in this interval, the next odd value is below X exactly when $\nu_2(3n + 1) \geq 2$: valuation one gives a value exceeding X , while valuation at least two gives a value at most X , with equality impossible for an odd output. Hence $n \equiv 1 \pmod{4}$. As $X \equiv 4 \pmod{36}$, the exact list satisfying these structural conditions is

$$n = X + 1 + 12j, \quad 0 \leq j \leq \frac{X - 4}{36}, \quad J_{\text{struct}} = \frac{X + 32}{36}. \quad (\text{C.structuralcount})$$

This counts candidates. Convergence and the remaining trajectory-height conditions still require certificates. Distinct seed descriptions cannot be counted twice if they give the same n .

There is a separate upper bound on the score obtainable within this geometry. For $J \geq 1$ distinct such complete blocks, put $K = K_q = 2 \cdot 3^{q-1}$ and $\ell = \lfloor K/(2J) \rfloor$. Remove the ℓ smallest-index roots from every block and exclude all their residue classes. At least $K/2$ residues remain, and every root at a remaining residue is at least $X4^\ell$. Assign total mass KP uniformly to the remaining residues. Its entries are at most $2P \leq 2 \leq B_q$, so it is feasible for (R.jointscore). Therefore

$$W_{\mathcal{F}}^{\text{joint}} \leq \frac{KPJ}{X} 4^{-\ell}. \quad (\text{C.geometryceiling})$$

At the canonical level, $N \geq 17$: $V_N \leq 1$ cannot hold at $N = 16$, since $t \geq 4b$, $\lambda > 1/4$, and $V_{16} > 2^{4b-32} > 1$. Consequently $q \geq 17b \geq 16b$, $K \geq 2^q \geq X^4$, and $J \leq J_{\text{struct}} \leq X$. It follows that

$$\ell \geq K/(4X) \geq 2^{12b-2}, \quad K \leq \ell^2 \leq 2^\ell, \quad W_{\mathcal{F}}^{\text{joint}} \leq 2^{-\ell} \leq 2^{-2^{402653182}}. \quad (\text{C.fixedceiling})$$

The floor costs at most a factor two, and $\ell \geq \sqrt{K}$, since $K \geq X^4 \geq 16X^2$. The integer inequality $\ell^2 \leq 2^\ell$ holds for $\ell \geq 4$. This bounds the stated mean-and-capacity allocation certificate. Since $c_{\text{comp}}(W) < W$ in the small-score domain, it also bounds that formula. It places no upper bound on the actual density of convergent integers. Other height geometries or stronger information about the reference marks are outside this counting ceiling.

Appendix D. Further consequences of the reference-law estimates

This appendix records consequences of the reference-law estimates beyond the three branches used in the explicit density recipe (D.densityrecipe). Convex-order inequalities propagate the whole concentration profile, whose finite allocation problem gives (D.profiledensity). Projectivity and summable mixing also give bounded relative entropy (F.entropybound). These results require no change to the root or counting constructions; a strict numerical improvement from the profile remains unproved.

The notation is inherited from Sections 9 and 10: f_n has mean one on G_n , H_m has mean $\mu = 8/9$, and $L = 2 \log 2$ is the limiting capacity coefficient. In Section D.2, W is a source-family score in the range of (D.mass), and $p_m = (W - \kappa \epsilon_m)_+$ is its coarse marked-mass bound. The symbols $D_s(n)$ and r_s in Section D.3 denote divergence and its rate; they are unrelated to the seed-prefix depth D in Appendix C.

D.1. Convex-order inequalities for the reference density and fan

The affine mixture (F.mixture), which proves the moment inequality (F.inheritance), also permits a general convex test.

Proposition D.1 (convex-order inequality). For every finite convex function $\Phi : [0, \infty) \rightarrow \mathbb{R}$ and integers $u, v \geq 0$,

$$\langle \Phi(f_{u+v}) \rangle_{u+v} \leq \frac{1}{3^{u+v}} \sum_{x \in G_u} \sum_{y \in G_v} \Phi(f_u(x)f_v(y)). \quad (\text{F.convex})$$

Proof. Fix y in (F.mixture). If $f_v(y) > 0$, divide its coefficients by their sum $f_v(y)$ and apply Jensen to the convex function $z \mapsto \Phi(f_v(y)z)$. Averaging over the affine permutations of G_u gives

$$3^{-u} \sum_{z \in G_u} \Phi(f_{u+v}(y + 3^v z)) \leq 3^{-u} \sum_{x \in G_u} \Phi(f_v(y)f_u(x)).$$

If $f_v(y) = 0$, both sides equal $\Phi(0)$. Average over y . The countable mixture can be grouped by the finitely many values of f_u , so ordinary finite Jensen suffices. A zero-length block gives equality. $\square$

Both compared distributions have mean one. Thus (F.convex) is a convex-order bound: every convex test has no larger average for the combined law than for the product of the two independent finite-table variables. Moment inheritance is its power-function specialization. The positive-part tests used in (D.hinge) will retain more of the upper tail than one power moment.

The same mixture gives the version needed for the whole fan:

$$\langle \Phi(H_{u+v}) \rangle_{u+v} \leq \frac{1}{3^{u+v}} \sum_{x \in G_u} \sum_{y \in G_v} \Phi(f_u(x)H_v(y)). \quad (\text{F.convexfan})$$

To prove this, split f_{u+v+1} into a prefix of length u and a suffix of length $v+1$. In the fiberwise proof of (F.convex), restrict the suffix residue to the 2 mod 3 coset and replace the test by $z \mapsto \Phi(4z/9)$. The bijections ϕ_{u+v} and ϕ_v turn the two sides into (F.convexfan). The suffix coset has 3^v elements, giving the displayed normalization. This also covers $v = 0$, for which the suffix has length one and $H_0 = 8/9$. Both sides have mean $8/9$. Repeated application therefore bounds convex tests by product distributions; it need not recover the exact distribution of H_{u+v} .

D.2. The full concentration profile and optimal residue allocation

The moment-to-mass estimate (D.generalp) is one way to use residue capacity. The following finite optimization records the strongest bound obtained from the whole marker distribution and that capacity alone.

Proposition D.2 (optimal capacity bound). Fix $m \geq 2$, put $Q = 3^m$, and let $0 \leq p \leq L\mu$. Define

$$\mathcal{C}_m(p) = \min \left\{ \sum_{r \in G_m} w_r : 0 \leq w_r \leq L/Q, \quad \sum_{r \in G_m} w_r H_m(r) \geq p \right\}. \quad (\text{D.allocation})$$

Then

$$\mathcal{C}_m(p) = \sup_{t > 0} \frac{[p - L\langle (H_m - t)_+ \rangle_m]_+}{t}, \quad 0 \leq \mathcal{C}_m(p) \leq p/\mu. \quad (\text{D.hinge})$$

Proof. Put $h_r = H_m(r)$ and $c = L/Q$. For every admissible allocation and $t > 0$,

$$p \leq \sum_r w_r h_r \leq t \sum_r w_r + c \sum_r (h_r - t)_+.$$

This proves the lower bound by the displayed supremum. To attain it when $p > 0$, order the entries as $h_1 \geq \dots \geq h_Q$ and let j be the least index with $c \sum_{i \leq j} h_i \geq p$. Feasibility follows from $c \sum_i h_i = L\mu$, and $h_j > 0$. Fill the first $j-1$ entries to capacity, allocate $(p - c \sum_{i < j} h_i)/h_j$ at entry j , and leave the others empty. The mass is

$$c(j-1) + \frac{p - c \sum_{i < j} h_i}{h_j}.$$

Choosing $t = h_j$ in (D.hinge) gives exactly this number, including ties. Thus there is an optimizer with at most one partially occupied entry. At $p = 0$, both sides are zero. The uniform allocation $w_r = p/(Q\mu)$ is feasible and has mass p/μ , proving the upper bound. $\square$

For any finite set $\mathcal{S} \subset (1, \infty)$ of orders with valid moment ceilings, the allocation optimum dominates all their capacity lower bounds simultaneously:

$$\max(\{u_m^{\text{comp}}(W)\} \cup \{v_{s,m}(W) : s \in \mathcal{S}\}) \leq \mathcal{C}_m(p_m) \leq W/\mu. \quad (\text{D.profiledominates})$$

The zero-residual case is immediate. Otherwise apply (D.capacity), (D.generalp), and (D.complementexplicit) to an allocation attaining $\mathcal{C}_m(p_m)$. Its mass is at most $p_m/\mu \leq W/\mu < L/2$, so the centered inverse uses exactly the increasing branch justified in (D.inverse). Hölder gives each other branch. There is no implied ordering of the bounds as s decreases. For example, for a constant marker one, capacity coefficient one and residual mass $1/2$, the order $2, 3/2, 5/4$ bounds are respectively $1/4, 1/8, 1/32$.

The feasible range is exact. At $p = L\mu$, all positive-marker entries may be filled and zero-marker entries left empty. If $p > L\mu$, no allocation is feasible, and the hinge supremum is infinite as $t \downarrow 0$. Optimality here concerns the residue-capacity constraints, not realization of an optimizing allocation by integer orbits.

For actual finite source families the coefficient is $T_X = T_{X,3^m,R_X}$, not L . If $p_X \geq 0$ is a guaranteed coarse marked mass at that cutoff, the correct finite inequality is

$$U_X \geq \frac{[p_X - T_X \langle (H_m - t)_+ \rangle_m]_+}{t}. \quad (\text{D.finitehinge})$$

For fixed m , any strict lower bound below $\mathcal{C}_m(p_m)$ can be witnessed by one fixed threshold $t > 0$. First choose that threshold, then a sufficiently small fixed graft tolerance, then a cutoff beyond which T_X and the coarse error bound have their required accuracy. This proves the analogue of (D.mass) with $\mathcal{C}_m(p_m)$ in place of $u_m(W)$. Its radial domain follows from $\mathcal{C}_m(p_m) \leq W/\mu$. Consequently (D.radial) also gives the valid fixed-level bound

$$\underline{d}(\mathcal{G}_c) \geq \mathcal{D}(\mathcal{C}_m(p_m)). \quad (\text{D.profiledensity})$$

Only a fixed threshold is needed before the cutoff limit.

The functions $z \mapsto (z - t)_+$ are increasing and convex. Thus (F.convexfan), and finite pointwise upper tables with directed rounding, can supply upper bounds for their averages. General convex tests need not be increasing and cannot automatically use such upper tables. An efficient representation of these propagated profiles would permit quantitative comparison with the explicit bound c_{comp} and the exact two-sided inverse in (D.twosidedcapacity).

D.3. Rényi divergence rates and bounded relative entropy

For $s > 1$ define the Rényi divergence from the uniform law, with relative entropy at order one, by

$$D_s(n) = \frac{\log M_s(n)}{s-1}, \quad D_1(n) = \langle f_n \log f_n \rangle_n, \quad 0 \log 0 := 0.$$

These are normalized using the mean-one density f_n .

Proposition D.3 (rates justified by inheritance). For every $s \geq 1$,

$$r_s := \lim_{n \rightarrow \infty} \frac{D_s(n)}{n} = \inf_{n \geq 1} \frac{D_s(n)}{n} \geq 0. \quad (\text{F.divergencerate})$$

The rates are nondecreasing in s , and $\lim_{s \downarrow 1} r_s = r_1$. For a fixed $s > 1$, if a block length $\ell \geq 1$ has certificates $M_s(j) \leq B_j$ for $0 \leq j \leq \ell$, with $B_0 = 1$, then

$$D_s(n) \leq n\bar{r}_s + O_{s,\ell}(1), \quad \bar{r}_s = \frac{\log B_\ell}{(s-1)\ell} \geq r_s. \quad (\text{F.blockrate})$$

The implicit constant here also depends on the fixed finite certificates.

Proof. For $s > 1$, (F.inheritance) makes D_s subadditive; Jensen gives nonnegativity. Finite-level differentiation at order one gives $D_s(n) \rightarrow D_1(n)$, including zero entries by continuity. Taking this limit in subadditivity proves it also for D_1 . For any nonnegative subadditive sequence a_n , write $n = k\ell + j$, $0 \leq j < \ell$. Then $a_n \leq ka_\ell + a_j$, so $\limsup a_n/n \leq a_\ell/\ell$. Taking the infimum over ℓ and using the opposite bound from the definition of the infimum proves (F.divergencerate). The same division into blocks using $\log B_j/(s-1)$ proves (F.blockrate).

For a finite mean-one density, Hölder makes $s \mapsto \log \langle f^s \rangle$ convex. Its value at one is zero, so its secant slope from one, D_s , is nondecreasing. Thus r_s is nondecreasing, and the two infima commute:

$$\inf_{s>1} r_s = \inf_{s>1} \inf_{n \geq 1} \frac{D_s(n)}{n} = \inf_{n \geq 1} \inf_{s>1} \frac{D_s(n)}{n} = r_1.$$

This proves the order-one limit. $\square$

For a fixed $s > 1$, subadditivity gives $D_s(n) = nr_s + o(n)$; by itself it does not upgrade that remainder to $O(1)$ at the true rate. The block rate $\bar{r}_s$ is a certified upper bound. At order one, projectivity and the summable mixing envelope give more.

Proposition D.4 (bounded order-one divergence). For every $n \geq 2$,

$$0 \leq D_1(n) - D_1(n-1) \leq \min\{\log 3, \epsilon_{n-1}\}. \quad (\text{F.entropyincrement})$$

Consequently $D_1(n)$ increases to a finite limit D_1^∞ , and

$$\begin{aligned} \sup_{n \geq 0} D_1(n) &= D_1^\infty \leq \frac{2}{3} \log 2 + \sum_{j \geq 1} \min\{\log 3, C_6 j^{-2314/25}\} < \infty, \\ 0 \leq D_1^\infty - D_1(n) &\leq \sum_{j \geq n} \min\{\log 3, \epsilon_j\} \quad (n \geq 1), \quad (\text{F.entropybound}) \\ r_1 &= 0, \quad \lim_{s \downarrow 1} r_s = 0. \end{aligned}$$

Proof. Fix $n \geq 2$ and put $g = f_{n-1} \circ \pi_{n,n-1}$. By the projectivity following (T.law), g is the average of f_n on each three-point fiber. Nonnegativity gives $f_n = 0$ on a zero- g fiber, and $0 \leq f_n/g \leq 3$ on every other fiber.

For $0 \leq u \leq 3$, define $\phi(u) = u \log u - u + 1$, with $0 \log 0 = 0$. Convexity and its minimum at one give $\phi(u) \geq 0$. If $u \leq 1$, then $u \log u \leq 0$, so $\phi(u) \leq 1 - u$. If $1 \leq u \leq 3$, its chord from one to three gives

$$\phi(u) \leq \frac{3 \log 3 - 2}{2} (u - 1) \leq u - 1.$$

Here $\log 3 = (\log_2 3) \log 2 < (8/5)(7/10) = 28/25 < 4/3$, using the elementary logarithmic inequalities already established in the construction. Thus $0 \leq \phi(u) \leq |u - 1|$ on the whole interval.

On each nonzero fiber, the sum of $f_n - g$ is zero and $\log g$ is constant. These facts cancel the linear and $\log g$ terms in the entropy difference, giving

$$D_1(n) - D_1(n-1) = \langle g \phi(f_n/g) \rangle_n \leq \langle |f_n - g| \rangle_n \leq \epsilon_{n-1}.$$

The integrand is defined as zero on zero- g fibers. Its nonnegativity proves the lower bound; the last inequality is precisely the unhalved bound (T.mix). The same fiber cancellation also writes the increment as $\langle f_n \log(f_n/g) \rangle_n$. Since $f_n/g \leq 3$ wherever $f_n > 0$ and $\langle f_n \rangle_n = 1$, this is at most $\log 3$, proving (F.entropyincrement).

At level zero $f_0 = 1$, while the even and odd geometric letters give $f_1(0) = 0, f_1(1) = 1, f_1(2) = 2$. Hence $D_1(0) = 0$ and $D_1(1) = (2/3) \log 2$. Summing (F.entropyincrement) and using (M.envelope) proves boundedness, since $2314/25 > 1$. Summing from $n+1$ onward gives the stated tail. Boundedness gives $r_1 = 0$, and (F.divergencerate) gives the right-hand limit in the order s . $\square$

For scale, the following values are certified to round to four decimal places; the complete record covers every $0 \leq n \leq 16$.

n	1	2	4	8	12	16
$D_1(n)$	0.4621	0.6304	0.8240	1.0087	1.1008	1.1544
$D_1(n)/n$	0.4621	0.3152	0.2060	0.1261	0.0917	0.0721

These are finite-level values, not an estimate of D_1^∞ . The comparison $\bar{r}_{3/2} = (\log b_{12})/6 \approx 0.1247$ is a certified block-rate upper bound from (F.blockrate), not an identified value of $r_{3/2}$. The proof identifies $r_1 = 0$.

The computation uses positive upper arrays as in Section 9, rather than treating their rounded entries as the exact law. For an upper density $h \geq f_n$, put $\eta = \langle h \rangle_n - 1$, $B = \max h$, and $T = \langle h \log h \rangle_n$. For $0 < \eta < 1$,

$$T - \eta(\log B + 1) \leq D_1(n) \leq T - \eta \log \eta. \quad (\text{F.entropyrounding})$$

Indeed, with $d = h - f_n \geq 0$, the identity $h \log h - f_n \log f_n = d \log h + f_n \log(h/f_n)$ and $\log(1+x) \leq x$ bound this difference above by $d(\log B + 1)$ and below by $d \log d$. Zero entries are interpreted by continuity. Averaging and convexity of $d \log d$ prove the two bounds; at $\eta = 0$, $h = f_n$. Rational logarithm enclosures and this correction give intervals of width less than 1.3×10^{-7} for the displayed entropies. The archive contains the rational interval endpoints determining the displayed rounding.

The tail estimate concerns convergence as the level n increases. It gives no quantitative rate as $s \downarrow 1$, no conclusion $r_s = 0$ at a fixed $s > 1$, and no polynomial moment bound at orders $3/2$ or 2 . Nor does it improve the optimized density constant without a further quantitative capacity comparison. For the fan, whose mean is $\mu = 8/9$, divergence is computed from H_m/μ . The capacity branches retain their residual-mass powers and capacity coefficient.

Appendix E. Primitive decay from the original pair process

We prove an explicit primitive Fourier-decay estimate for the reference law. The argument uses the pairing and separated-triangle method of Tao [1, Section 7] and its quantitative development in [3]. Working directly with the original pair process gives an exact binomial first-passage law. Its exponential moment yields a linear vertical window and the degree-four recurrence used to specify the coefficient.

E.1. Primitive Fourier decay and its coefficient

Let independent positive integers K_i satisfy $\Pr(K_i = a) = 2^{-a}$, and put

$$X_n = \sum_{i=1}^n 3^{i-1} 2^{-(K_1 + \dots + K_i)} \pmod{3^n}, \quad F_n(\xi) = \mathbb{E} e^{2\pi i \xi X_n / 3^n}. \quad (\text{E.law})$$

This is (T.law), with the unnormalized positive Fourier convention of Section 6. All inverse powers of two are taken in the indicated residue ring.

For a positive integer B , define

$$\epsilon = 2^{-78}, \quad d = \epsilon^2/2, \quad \eta = d/(32B), \quad \delta = (16 \cdot 10^B)^{-1}, \\ K = 16(B+1)/\epsilon^2, \quad R = 4(K + 16(B+1)) + 1, \quad A = 64/\delta, \quad H = 2^{17} A^2/\delta.$$

The symbols K, R, A, H in this appendix are local parameters; in particular K is not a valuation letter K_i . They are positive integers. Set

$$u(p) = H(p+1)^4, \quad v_0 = K, \quad v_{i+1} = v_i + u(v_i) + K + 2 \quad (0 \leq i < R), \\ P = v_R + K + 2, \quad Z = 128A(P+1)^2/\delta, \\ M = \lceil \max \{64B/d, 20(P+1), (2560(B+1))^2, Z^2/\eta, 4096\} \rceil, \quad C_B^{\text{loc}} = (3M)^B. \quad (\text{E.recipe})$$

This is a finite integer recipe fixed before the level and frequency.

Theorem E.1 (local primitive decay). For every positive integer B , every $n \geq 1$, and every ξ with $3 \nmid \xi$,

$$|F_n(\xi)| \leq C_B^{\text{loc}} n^{-B}. \quad (\text{E.primitive})$$

For real requested order $b > 0$, the integer order $B = \lceil b \rceil$ supplies the same assertion with n^{-b} . We prove the theorem below.

E.2. Pairing and the discounted process

Fix n, ξ as in the theorem and let $N = \lfloor n/2 \rfloor$. The case $n = 1$ follows from $|F_n| \leq 1$; assume $n \geq 2$. Put $B_j = K_{2j-1} + K_{2j}$ and $S_j = B_1 + \dots + B_j$, with $S_0 = 0$. Then

$$\Pr(B_j = b) = (b-1)2^{-b} \quad (b \geq 2).$$

Conditional on a pair sum b , its first letter is uniform on $1, \dots, b-1$. The contribution of that pair, after j preceding pairs and total valuation l , is

$$3^{2j} 2^{-l} (2^{-a} + 3 \cdot 2^{-b}).$$

The conditional character is a unit-modulus factor times $(b-1)^{-1} \sum_{a=1}^{b-1} e_{3^n}(\xi 3^{2j} 2^{-(l+a)})$. Its modulus is at most one. When $b = 3$, that modulus is $\cos(\pi\theta(j, l))$, where

$$\theta(j, l) = \text{bal} \left(\frac{\xi 3^{2j} 2^{-(l+2)}}{3^n} \right) \in (-1/2, 1/2] \quad (0 \leq j < N, l \in \mathbb{Z}). \quad (\text{E.phase})$$

Here bal selects the balanced representative modulo one. Call a point black if $|\theta| \leq \epsilon$, and white otherwise. Since $1 - \cos(\pi t) \geq 2t^2$ for $|t| \leq 1/2$, the pair factor at a white state is at most $1 - 2\epsilon^2$ when the next pair sum is three.

Given the pair sums, different first-letter choices are independent, and the state of each pair depends only on earlier pair sums. Multiplying these conditional bounds and bounding a possible last odd letter by one gives

$$|F_n(\xi)| \leq U(0, 0),$$

where $U(j, l)$ is the expectation, from that state, of the product up to column N of

$$(1 - 2\epsilon^2)^{\mathbf{1}_{W(j+t, l+S_t)} \mathbf{1}_{\{B_{t+1}=3\}}}. \quad (\text{E.kill})$$

Set $U = 1$ at and beyond N ; in particular $0 \leq U \leq 1$. At a white state the one-step kernel has total mass $1 - \frac{1}{4}(2\epsilon^2) = 1 - d$. For products of the discount factors we use $(1 - 2\epsilon^2)^k \leq e^{-2\epsilon^2 k}$.

E.3. Black triangles and separation

Lemma E.2 (geometry). The black lattice points are partitioned into triangles

$$\Delta = \{(j, l) : j \geq j_0, l \leq l_0, (j - j_0) \log 9 + (l_0 - l) \log 2 \leq t\}, \quad (\text{E.triangle})$$

with integer corners (j_0, l_0) , $j_0 \geq 0$, and $t \geq 0$. Every triangle lies in $j \leq N - 16$. Every lattice point in the state space $0 \leq j < N$, outside a triangle but within distance sixteen of it, is white. Here distance is Euclidean distance to the triangle's lattice points. Distinct triangles have disjoint lattice points.

Proof. The exact phase relations are

$$\theta(j+1, l) = 9\theta(j, l) \pmod{1}, \quad \theta(j, l-1) = 2\theta(j, l) \pmod{1}. \quad (\text{E.phase-steps})$$

Call a point weakly black if its phase has modulus at most $1/100$. Three elementary implications will be used. A weakly black point with a black east or south neighbor is black, since multiplication by nine or two does not wrap at that size. If east and south neighbors are weakly black, the identity $\theta = \theta_{\text{east}} - 4\theta_{\text{south}} \pmod{1}$ first bounds the parent by $5/100$; its doubled phase then equals the south phase as a real number and bounds it by $1/200$. If west and south neighbors are weakly black, the west relation first bounds the parent by $9/100$; doubling again bounds it by $1/200$.

Starting from a black point (j, l) , move up along its black run to the last black point (j, l_0) , then left along the black row to its first black point (j_0, l_0) , possibly at $j_0 = 0$. The upward run is finite: otherwise its nonzero phase would halve indefinitely, contradicting its minimum modulus $3^{-(n-2j)}$. Write $|\theta(j_0, l_0)| = \epsilon e^{-t}$. Multiplication along the known black row and column does not wrap, so the starting point belongs to (E.triangle). The same multiplication formula makes every lattice point in that triangle black.

For the collar put $r = 16$. All amplifications below are at most 18^{r+1} , and $100 \cdot 18^{17} < 2^{78}$. Thus they carry a phase of modulus at most ϵ to one of modulus less than $1/100$. For a point southeast of the corner and within distance r , its excess logarithmic weight over the triangle is at most $r(\log 9 + \log 2)$. Outside the triangle its phase is therefore strictly between ϵ and $1/100$, with no wraparound; it is white.

For a northeast point (j', l') with $l_0 < l' \leq l_0 + r$, suppose it were black. Multiply down to height $l_0 + 1$. Propagate weak blackness left to column j using east and south neighbors if $j' \geq j$, or right using west and south neighbors if $j' < j$. The necessary lower row inside the triangle is black. Beyond its last top-row column it is still weakly black: at most sixteen additional east steps give modulus at most $\epsilon 9^{16} < 2^{-27} < 1/100$. The resulting weakly black point $(j, l_0 + 1)$, with black south neighbor, is black, contradicting the choice of l_0 .

For a point west of the corner, $j_0 - r \leq j' < j_0$, suppose it were black. Amplify east to column $j_0 - 1$. If $l' \geq l_0$, amplify down to l_0 ; if $l' < l_0$, propagate weak blackness upward using east and south neighbors. The east column is weakly black throughout the required range because $(l_0 - l') \log 2 \leq t + r \log 2$. The point $(j_0 - 1, l_0)$ is weakly black and has a black east neighbor, so is black, contradicting the choice of j_0 . There is no west case when $j_0 = 0$. These cases prove the collar.

Each lattice triangle is connected to its corner by north and west unit steps. Two intersecting lattice triangles must have the same lattice set: a path from an intersection to a point in their difference would cross the white collar. They then have the same extreme corner and, by its phase, the same t . This gives the asserted partition. Finally the corner phase is at least $3^{-(n-2j_0)}$, whence

$$j_0 + t/\log 9 \leq n/2 + \log \epsilon / \log 9.$$

Since $9^{17} < 2^{78}$, this places the whole triangle before $j = N - 16$. $\square$

E.4. First passage and white exits

For integer $s \geq 0$, let $J_s = \min\{j \geq 1 : S_j > s\}$ and $H_s = S_{J_s} - s$. Realize each geometric letter as the waiting time for one success of fair coin tosses. If T_s counts the successes in the first s tosses, then

$$T_s \sim \text{Bin}(s, 1/2), \quad J_s = \lfloor T_s/2 \rfloor + 1. \quad (\text{E.passage})$$

Conditional on T_s , the overshoot waits for two more successes when T_s is even and one when it is odd. Future tosses are independent. Consequently, for integers $j, h \geq 1$, with out-of-range binomial coefficients zero,

$$\Pr(J_s = j, H_s = h) = 2^{-s-h} \left[\binom{s}{2j-2} (h-1) + \binom{s}{2j-1} \right]. \quad (\text{E.joint})$$

In particular

$$J_s \leq s/2 + 1, \quad \mathbb{E}H_s \leq 4, \quad \Pr(H_s > h) \leq (h+1)2^{-h}. \quad (\text{E.overshoot})$$

For $s \geq 1, a > 0$, the variance $s/4$ of T_s gives

$$\Pr(|J_s - s/4| > a\sqrt{s} + 1) \leq (16a^2)^{-1}, \quad \sup_j \Pr(J_s = j) \leq 4/\sqrt{s}. \quad (\text{E.passage-bounds})$$

For the atom bound, each atom is the sum of two binomial atoms. The central estimate $\binom{2k}{k}/4^k \leq 1/\sqrt{k+1}$ follows by induction: its squared step is $(2k+1)^2(k+2) \leq 4(k+1)^3$, with difference $3k+2$. The odd case follows from the neighboring even case, so one binomial atom is at most $\sqrt{2/(s+1)}$, proving the displayed bound.

For $s \geq 1$ and $t \geq 0$, we also use

$$\Pr(T_s - s/2 \geq t) \leq e^{-2t^2/s}. \quad (\text{E.binomial-tail})$$

Indeed its centered exponential moment is $\cosh(a/2)^s \leq e^{sa^2/8}$, by integrating $\tanh x \leq x$; exponential Markov with $a = 4t/s$ gives the claim.

Start at (j, l) in a triangle with top l_0 , and set $s = l_0 - l$. Except with probability e^{-8} , $J_s \leq s/4 + \sqrt{s} + 1$; for $s = 0$ this holds deterministically. The top right corner is at least $s \log 2 / \log 9$ columns to the right of j . Since $\log 2 / \log 9 > 5/16$ and $\sqrt{s} \leq s/16 + 4$, the exit's horizontal overshoot past the last top-row lattice point is less than six. Also $H_s \leq 8$ except with probability $9/256$. On their intersection the exit lies outside the triangle within distance ten, and is before N . Its collar makes it white. Thus

$$\Pr((j + J_s, l_0 + H_s) \text{ is white}) \geq 1 - e^{-8} - 9/256 > 15/16. \quad (\text{E.white-exit})$$

E.5. Long crossings and the degree-four threshold

Observe the process $p \geq 0$ steps after that exit. Its position is $(j + J_s + p, l_0 + V_p)$, where $V_p = H_s + B_{J_s+1} + \dots + B_{J_s+p}$. For $1 < z < 2$, a geometric letter has moment $z/(2-z)$. Conditional on the parity at time s , V_p is the sum of either $2p+1$ or $2p+2$ fresh geometric letters. Hence $\mathbb{E}(4/3)^{V_p} \leq 4^{p+1}$. With $Y = A(p+1)$, using $\log(4/3) \geq 1/4$, $\log 4 < 2$, $A \geq 32$ and $e^x \geq x^2/2$, we obtain

$$\Pr(V_p > Y) \leq e^{-(A/4-2)(p+1)} \leq e^{-A(p+1)/8} \leq \frac{128}{A^2(p+1)^2} \leq \frac{4}{A(p+1)^2}. \quad (\text{E.vertical-tail})$$

Together with (E.passage-bounds), this puts the point in the rectangle

$$|x - (j + s/4 + p)| \leq W := A(p+1)\sqrt{s} + 1, \quad l_0 < y \leq l_0 + Y,$$

except with probability $1/(16A^2(p+1)^2) + 4/(A(p+1)^2)$. No conditioning on the realized vertical coordinate is used.

Assume $\sqrt{s} \geq Z$ and $p < P$. The recipe implies $W, Y, p \leq s/64$, $Y \leq W$ and $W \leq 2A(p+1)\sqrt{s}$. The horizontal range of the rectangle, even with Y subtracted at its left end, lies in the old triangle's top segment: use $\log 2 / \log 9 - 1/4 > 1/16$ for its right end.

Consider a different triangle of size $t' \geq u = u(p)$ meeting the rectangle at a lattice point (x, y) , with left column a . Then $x - Y < a \leq x$. Otherwise $(x - Y, l_0)$ would lie in both triangles: moving left by Y saves $Y \log 9$ of weight and moving down costs at most $Y \log 2$. Also a is in the old top segment. Let b' be the lower endpoint of the new triangle's real vertical section at column a . Disjointness at the lattice point (a, l_0) gives $b' > l_0$. Moving from (x, y) to (a, y) decreases the triangle weight, so $b' \leq y \leq l_0 + Y$.

At the common integer height $l_* = l_0 + \lceil u/(2 \log 2) \rceil$, each such triangle has a horizontal lattice interval starting at a , of length at least $u/12$. Its top, at $b' + t'/\log 2$, is above l_* . Since $u \geq 4Y \log 2$ and $u \geq 64$, the real section length is

$$\frac{(l_* - b') \log 2}{\log 9} \geq \frac{u/2 - Y \log 2}{\log 9} \geq \frac{u}{4 \log 9} > \frac{u}{10}.$$

Taking an integer part loses less than one, while $u/10 - 1 \geq u/12$. These lattice intervals are disjoint, so their left columns are separated by at least $u/32$.

All such left columns lie in an interval of length at most $2W + Y + 2$. Each makes at most $Y + 1$ columns of the rectangle bad, because $x - Y < a \leq x$. Thus the total number of bad columns is at most

$$(Y + 1)(1 + 32(2W + Y + 2)/u) \leq 2Y + 320YW/u.$$

This is a deterministic set of columns covering the whole rectangle. Apply the unconditional atom bound in (E.passage-bounds), then add the rectangle error. Since $u(p) = H(p+1)^4$,

$$\Pr(\text{triangle size at } J_s + p \geq u(p)) \leq \frac{4/A + 1/(16A^2) + 2560A^2/H}{(p+1)^2} + \frac{8A(p+1)}{\sqrt{s}}. \quad (\text{E.large-triangle})$$

Use $\sum_{r \geq 1} r^{-2} \leq 2$, $\sum_{r=1}^P r \leq (P+1)^2/2$ and $\sqrt{s} \geq Z$ to get

$$\Pr(\exists p < P : \text{triangle size at } J_s + p \geq u(p)) \leq 2(4/A + 1/(16A^2) + 2560A^2/H) + \delta/32 < \delta. \quad (\text{E.large-union})$$

E.6. White visits and predictable cancellation

For this subsection extend all columns $j \geq N$ as white. Write the walk as $z_t = (j + t, l + S_t)$. Define its first entry by $\sigma_1 = \inf\{t \geq 0 : z_t \text{ is black}\}$. At an entry $\sigma_i < \infty$, let l_i be the top of its triangle and put

$$\tau_i = \inf\{t > \sigma_i : l + S_t > l_i\}, \quad \sigma_{i+1} = \inf\{t \geq \tau_i : z_t \text{ is black}\},$$

with $\inf \emptyset = \infty$. In particular, a black exit is the next entry at that same time. The top l_i stays fixed during its crossing. The gaps before the first entry and between each exit and the next entry are white, and the exit times are strictly increasing.

By (E.white-exit) and the strong Markov property, the conditional expectation of $e^{-1_{\{\text{white exit}\}}}$ at each entry is at most $1 - (15/16)(1 - e^{-1}) \leq 3/4$. Iterating with the indicator that the next entry exists gives

$$\Pr(R \text{ entries before } P, \# \text{white states before } P < K) \leq e^K (3/4)^{R-1} \leq \delta. \quad (\text{E.entries})$$

The iteration is valid even if entries cease: that path contributes zero to the event of R entries. On the counted event, the first $R - 1$ exit factors concern times before P , and their white exits are a subset of all white states. The final bound follows from $\log(4/3) \geq 1/4$, so its logarithm is at most $-16(B+1)$, whereas $\log(1/\delta) < 4 + 3B$.

Now start immediately after the long crossing in E.5. Outside the event (E.large-union), every triangle entered at time $p < P$ has size less than $u(p)$. Its top is at distance less than $u(p)/\log 2 < 2u(p)$, so is passed within $u(p) + 1$ pair steps. If fewer than K white states occur, the initial gap and each later gap have length at most K . The recurrence defining v_i bounds successive entry and exit times. Fewer than R entries with fewer than K whites cannot fill $P = v_R + K + 2$ steps. The simultaneous quantifier over all $p < P$ in (E.large-union) includes the random entry times without another union bound. Consequently

$$\Pr(\# \text{white states in these } P \text{ steps} < K) \leq 2\delta.$$

At each white state, the next pair sum is sampled freshly and equals three with probability $1/4$. The state is known before that sample. The marks at the first K white states are therefore independent Bernoulli $(1/4)$ variables; the white extension makes these sampling times finite. Their lower-tail bound is

$$\Pr(\text{fewer than } K/8 \text{ marks among the first } K \text{ whites}) \leq e^{-K/32} \leq \delta.$$

For completeness, the centered log moment of a Bernoulli variable has second derivative at most $1/4$; integrating twice gives an upper bound $a^2/8$. Exponential Markov at deviation $K/8$ gives the stated exponent. On $K/8$ marked whites, the product in (E.kill) is at most $e^{-\epsilon^2 K/4} = e^{-4(B+1)} \leq \delta$. We conclude that the auxiliary product over the P post-exit steps obeys

$$\mathbb{E} \prod_{p=0}^{P-1} (1 - 2\epsilon^2)^{\mathbf{1}_{W(z_{J_s+p})} \mathbf{1}_{\{B_{J_s+p+1}=3\}}} \leq 4\delta = \frac{1}{4 \cdot 10^B}. \quad (\text{E.postexit})$$

E.7. Weighted induction and the terminal event

The function $U(j, l)$ is periodic in l , with period dividing $2 \cdot 3^{n-1}$. For $0 \leq m \leq N$, set

$$D_m = \max_{N-m \leq j \leq N, l \bmod 2 \cdot 3^{n-1}} \max(N-j, 1)^B U(j, l).$$

Then $D_0 = 1$, and $D_m \geq D_{m-1}$. We prove the reverse inequality when $M \leq m \leq N$. Only the new layer $j = N - m$ needs checking. Put $D = D_{m-1} \geq 1$; it bounds all strictly later states with their displayed weights.

At a white state with remaining distance $r \geq 8B/d$ whose later states are controlled by D , the kernel mass $1 - d$ gives

$$U(j', l') \leq (1 - d)D(r - 1)^{-B} \leq (1 - d/2)Dr^{-B}. \quad (\text{E.white-bound})$$

Indeed $\log(r/(r - 1)) \leq 2/r$ and $e^{d/4} \leq 1 + d/2$. This handles a white state on the new layer.

For a black state let s be its vertical distance to the triangle's top. If $s \leq \eta m$, stop at J_s . Deterministically,

$$J_s/m \leq d/(32B), \quad m - J_s \geq m/2 \geq 8B/d, \quad (1 - J_s/m)^{-B} \leq 1 + d/8.$$

Here $J_s \leq \eta m/2 + 1$ and $m \geq 64B/d$. This includes $s = 0$, where $J_0 = 1$. Discard all factors before the stop. At white exits use (E.white-bound), and elsewhere use the ordinary later-layer bound. By (E.white-exit),

$$U(j, l) \leq Dm^{-B}(1 + d/8)(1 - 15d/32) < Dm^{-B}. \quad (\text{E.short})$$

If $s > \eta m$, the recipe gives $\sqrt{s} \geq Z$. The triangle fits before N , so $s \log 2 \leq m \log 9$ and $s < 16m/5$. Also $P \leq m/20$. Thus $J_s + P > 9m/10$ implies $J_s > 17m/20$, and (E.passage) gives $T_s - s/2 > m/10 - 2 \geq m/20$. The binomial tail implies

$$\Pr(J_s + P > 9m/10) \leq e^{-m/640}. \quad (\text{E.terminal-tail})$$

On the complementary event all sampled states are before N , and the remaining distance is at least $m/10$. Discard factors before J_s , use (E.postexit), and then the later-layer bound. On the exceptional event use one. Since $D \geq 1$,

$$\frac{U(j, l)}{Dm^{-B}} \leq 10^B \mathbb{E}[\text{post-exit auxiliary product}] + m^B e^{-m/640} \leq \frac{1}{4} + \frac{1}{4} < 1. \quad (\text{E.long})$$

For the second term, $\log m \leq \sqrt{m}$ and $m \geq (2560(B + 1))^2$ give $m^B e^{-m/640} \leq e^{-m/1280} \leq 1/4$. On the good event, every factor of the auxiliary product is before column N and agrees with the original product. On the exceptional event, the bound uses $U \leq 1$ and its probability in (E.terminal-tail).

This proves $D_m = D_{m-1}$ for $m \geq M$. For $m < M$, $U \leq 1$ gives $D_m \leq M^B$. Hence $U(0, 0) \leq M^B \max(N, 1)^{-B}$. Since $\lfloor n/2 \rfloor \geq n/3$ for $n \geq 2$, (E.primitive) follows with $C_B^{\text{loc}} = (3M)^B$. This also completes the odd-length case. $\square$

E.8. Size of the coefficient

For $B = 6409$, the scalar definitions give

$$K < 2^{173}, \quad R < 2^{175}, \quad A < 2^{25646}, \quad \delta^{-1} < 2^{25640}, \quad H + 3 < 2^{77000}.$$

Writing $w_i = v_i + 1$, we have $w_{i+1} \leq (H + 3)w_i^4$ and $w_0 < 2^{174}$. Therefore

$$\log_2 w_R < (174 + 77000/3)4^R < 2^{15}4^R, \quad \log_2(P + 1) < 2^{16}4^R.$$

Also $Z < 2^{51293}(P+1)^2$, $\eta^{-1} < 2^{175}$, and $M < 2^{102862}(P+1)^4$. The deliberately loose last prefactor bounds every entry in (E.recipe). It follows that

$$\log_2 M < 2^{19}4^R, \quad \log_2 C_{6409}^{\text{loc}} < 2^{33}4^R, \quad \log_2 \log_2 C_{6409}^{\text{loc}} < 33 + 2R < 2^{176}. \quad (\text{E.local-size})$$

For comparison, write the coefficient in [3] as $C_{\text{primitive}}^{\text{Maz}}$. Its finite recipe is

$$\begin{aligned} B_0 &= 6409, \quad E_0 = 2^{170}, \quad L_0 = 2^{80}, \quad T_0 = 10B_0 2^{3E_0}, \quad R_0 = 2^{E_0}(T_0 + 3 \cdot 8195 + 1) + 1, \\ p_0 &= T_0, \quad p_{i+1} = p_i + 10 \cdot 8^{32768}(p_i + 1)^3 + 1 + T_0, \quad P_{\max} = p_{R_0-1}, \quad P' = P_{\max} + 1, \quad S_0 = 2^{34}P'. \end{aligned} \quad (\text{M.primitiveiteration})$$

Set

$$\begin{aligned} T_1 &= (128B_0 2^{3E_0})^2 + 2, \quad T_2 = \max\{2(T_1 + L_0^2 + 1) + 4, 2^{8192B_0 2^{3E_0}}, 2\}, \\ T_{3,s} &= 2 + P' + P_{\max}^{10} + (8^{32768}(1 + P_{\max})^3)^3 + (100S_0)^2, \\ T_{3,t} &= (2^{38}B_0)^2 + 2^{38}(P' + 65), \quad T_* = \max\{T_1, T_2, 1, T_{3,s}, T_{3,t}\}, \quad C_{\text{primitive}}^{\text{Maz}} = (32B_0 T_*)^{B_0}. \end{aligned} \quad (\text{E.mazur-coefficient})$$

This is the coefficient recipe in [3], used for the numerical comparison below. Put $r = R_0 - 1$. Then $r > 2^{4E_0}$, $p_0 \geq 2$, and $p_{i+1} \geq p_i^3$, so $p_r \geq 2^{3^r}$. The coefficient is at least p_r through its $P_{\max}^{10}$ term. With $L_3(x) = \log_2 \log_2 \log_2 x$, we obtain

$$L_3(C_{6409}^{\text{loc}}) < 176, \quad L_3(C_{\text{primitive}}^{\text{Maz}}) > 2^{172}, \quad C_{6409}^{\text{loc}} < C_{\text{primitive}}^{\text{Maz}}. \quad (\text{E.coefficientcomparison})$$

The comparison follows from the defining recurrences and their growth bounds, without expansion of the resulting integers.

E.9. Mixing exponents and coefficient dependence

For a primitive input of order β , the two leading powers of the fixed-width conversion are

$$a_\beta(v) = \beta - 2 - \frac{1}{2}(\log 2)v^2, \quad q(v) - 1 = v \log 2 - 1. \quad (\text{E.budget})$$

For $\beta > A_* + 2$, both exceed A_* precisely on the intersection of $[80, 679/5]$ with

$$((A_* + 1)/\log 2, \sqrt{2(\beta - 2 - A_*)/\log 2}).$$

For $\beta \leq A_* + 2$, the intersection is empty. At width 80 the smallest integral β making both exceed six is

$$\beta \geq 2227. \quad (\text{E.floor})$$

At $\beta = 2227$, the smaller power is about 6.929. The order-92 seed estimate and geometric level contraction use the stronger choice $\beta = 6409$: the rational width 6749/50 gives both powers strictly above 2314/25, as in (M.exponentguards). Thus the entire Section 6 envelope and order-92 seed estimate use the new coefficient without a change of exponent.

Lower-bound guards and quantitative sensitivity have different roles. For any valid primitive coefficient $C \geq 1$, the adapter $C_6(C) = 2C20^{6409} + 2$ satisfies $C_6(C) > 2^{27700}$, enough for (B.coefficientscope), $S_{92} < C_6$, and the small-level mixing bound. It remains strictly increasing in C . In particular

$$F_b(C) = 2^{467}b16^b(C_6(C) + 1), \quad V_C(n) = 2^{\lceil \log_2 F_b(C) \rceil} T_3(n, \lambda)$$

are coefficient-dependent, and the envelope decreases when C decreases. The factor 20^{6409} does not erase a decrease of C . The coefficient comparison can therefore improve certified formulas; Appendix B traces this dependence through startup, allocation, and the density formula. The clock threshold remains c_0 .

Formalization and computational data

Scope of the Lean formalization

A companion Lean 4 development supplies eighteen selected statements. The main convergence theorem gives one positive lower natural-density constant for every clock strictly above the threshold; the fixed-target theorem gives one such constant for each positive target not divisible by three. These are existence statements: the formal release does not claim the complete optimized density formulas of this paper.

The selected results have the following scope.

Result family	Formalized conclusion
Collatz convergence to one	Common-clock positive lower density, the rational clock specialization, and selected depth-eleven second-moment and fractional bounds.
Fixed-target hitting	Positive lower density for every positive target not divisible by three, the exact target criterion, and corresponding depth-eleven bounds.
Corollaries	A positive individual-root consequence and vanishing-clock-loss statements for convergence to one and for each admissible target.
Concentration profiles	Exact finite fan allocation and its hinge dual, each real-moment lower bound, and profile-density bounds at an existential positive paid level.
Reference entropy	Increment bounds, a finite monotone limit with an explicit tail, and zero normalized growth, including the positive-index normalized infimum.

The primitive Fourier-decay argument of Appendix E is proved within the formal development, including its all-order coefficient recipe and the order-6409 specialization used in Section 6. The selected proof dependencies do not assume Mazur’s numerical primitive-decay theorem. Adapted probability and word constructions retain their original attribution and license notices. The coefficient-size estimate, its comparison with the historical coefficient, and the resulting optimized-density comparisons remain paper proofs. The other divergence-rate results, the Rényi-order right limit, complementary capacity, deeper tables, full optimization, and amplified root families are outside the selected Lean release.

Source record and Palomar registration

The companion development defines the ordinary Collatz map, logarithmic hitting sets, and lower natural density directly. Its theorem map records the correspondence with this paper, and its dependencies are pinned to specified Lean and Mathlib versions. The eighteen selected statements have passed local compilation and axiom audits.

Public release through the [Palomar registry](#) is pending. No registry identifier or version has yet been assigned, and independent registry verification is not claimed in this preprint. The scope table above describes the companion development; the manuscript proofs stand independently of its eventual registration.

Computational data and reproducibility

The finite certificates are specified by the recurrences, pseudocode, and integer inequalities in Section 9 and the appendices. The reference-density tables use exact integer arithmetic, upward rounding, and explicit positive-tail bounds. The depth-sixteen cyclic calculation covers all $3^{16} = 43,046,721$ residues. Fractional-moment bounds use integer square roots, and entropy values use rational logarithm enclosures. Section 9 specifies the transfer recurrences and certificate inequalities. The supplementary computations also include scalar comparisons, root-trajectory certificates, and residue-counting calculations used in the explicit bounds; these broader paper computations should be distinguished from the selected formal proof package.

Use of AI-assisted tools

AI-assisted tools were used in manuscript preparation, Lean proof development, code preparation, and editorial revision.

References

1. Terence Tao, *Almost all orbits of the Collatz map attain almost bounded values*, Forum of Mathematics, Pi 10 (2022), e12. [Corrected author preprint, v7](#), 16 July 2026.
2. Terence Tao, *Equidistribution of Syracuse random variables and density of Collatz preimages*, 25 January 2020. [Author's exposition](#).
3. L. Mazur, *Explicit Positive-Density Collatz Convergence in Logarithmic Time*, manuscript v2.1, 6 September 2026. [Manuscript and supplementary material](#).
4. L. Mazur, *Positive Lower Density of Collatz Predecessors*, manuscript v1.0, 6 September 2026. [Manuscript and theorem statements](#).
5. Alex Kontorovich and Jeffrey C. Lagarias, *Stochastic Models for the $3x + 1$ and $5x + 1$ Problems*, in Jeffrey C. Lagarias (ed.), *The Ultimate Challenge: The $3x + 1$ Problem*, American Mathematical Society, 2010, pp. 131–188. [Author preprint](#).
6. Ilia Krasikov and Jeffrey C. Lagarias, *Bounds for the $3x + 1$ Problem using Difference Inequalities*, Acta Arithmetica 109 (2003), 237–258. [Author preprint](#); [published article](#).
7. Idris Ali Shaik, *Polylogarithmic Descent for Almost All Collatz Orbits in Natural Density*, preprint, 6 September 2026. [SSRN 7290240](#), [doi:10.2139/ssrn.7290240](#).
8. Idris Ali Shaik, *Quantitative Endpoint Transport for the Collatz Map with a Natural-Density Descent Application*, preprint, 6 September 2026. [SSRN 7341398](#), [doi:10.2139/ssrn.7341398](#).
9. Riho Terras, *A stopping time problem on the positive integers*, Acta Arithmetica 30 (1976), no. 3, 241–252. [doi:10.4064/aa-30-3-241-252](#).
10. C. J. Everett, *Iteration of the number-theoretic function $f(2n) = n$, $f(2n + 1) = 3n + 2$* , Advances in Mathematics 25 (1977), no. 1, 42–45. [doi:10.1016/0001-8708\(77\)90087-1](#).
11. Ivan Korec, *A density estimate for the $3x + 1$ problem*, Mathematica Slovaca 44 (1994), no. 1, 85–89. [DML-CZ article record](#).
12. Jeffrey C. Lagarias, *The $3x + 1$ problem and its generalizations*, American Mathematical Monthly 92 (1985), no. 1, 3–23. [Author's article](#).
13. Manuel Inselmann, *An approximation of the Collatz map and a lower bound for the average total stopping time*, preprint (2024), arXiv:2402.03276. [Author preprint](#).