ANNEALED EQUIDISTRIBUTION OF THE BINARY SUM OF DIGITS ALONG SYRACUSE ITERATES, TO DEPTH $c \log N$

LEWIS VALENTINE

Date: July 2026 — draft v6.

Abstract. Let $\mathrm{Syr}(n) = (3n+1)/2^{\nu_2(3n+1)}$ denote the Syracuse (accelerated Collatz) map on odd integers, m_k its k-th iterate, and $\mathrm{s}_2(n)$ the binary sum of digits (Hamming weight, popcount). We prove that for every $\alpha \notin \mathbb{Z}$ and every fixed k, $\sum_{n\le N,\ n \text{ odd}} \mathrm{e}(\alpha\, \mathrm{s}_2(m_k(n))) = o(N)$, an *annealed* (seed-averaged) equidistribution statement for the binary digits of Syracuse iterates; and that the same holds uniformly for depths $k \le (c_0 - \varepsilon)\log_2 N$ with the unconditional constant $c_0 \approx 0.124$. The bottleneck is a single-modulus equidistribution problem for s_2 along arithmetic progressions of difference 3^k (Problem T1), for which all known power-saving results are averages over the modulus. We give an exact orbit-decomposition of the associated second-moment transfer sum: the full-period products admit a closed cyclotomic evaluation (at the parity character the orbit product equals 9 at *every* 3-adic level, locating the Newman–Coquet resonance exactly and only at level one), and the deep-window regime is exactly diagonal. This reduces the expected mean-square bound, and with it a depth constant $c_0 = \frac{1}{2} - \varepsilon$ matching the information-theoretic ceiling, to one arithmetic inequality on signed binary representations of multiples of 3^m, together with a mass-weighted non-concentration bound for cylinder offsets (a count-multiplicity version stated in an earlier draft is *refuted* here, by pigeonhole in the fixed unit group modulo 3^k; the repaired hypothesis is on the ℓ^2 mass of the offset profile). That residual signed-weight input is here identified with a non-concentration property of the Thue–Morse spectral measure on 3-adic sublattices. On this path we now prove, unconditionally, three of the previously measured facts: an exact power lower bound $M(N) \ge \frac{2}{9}N^\beta$ with $\beta = \log_2 \frac{1+\sqrt{17}}{4}$ for the Thue–Morse mean-square (a transfer-operator eigenvalue), the exact finite spectrum of the associated transfer operator (spectral radius $9^{1/P_t}/2 \downarrow \frac{1}{2}$, so each fixed level contracts geometrically), and the doubling-invariance of the Thue–Morse measure and its autocorrelation; combining the first with a renormalization estimate yields an *unconditional* per-level contraction $L(n,N) \le C\, M(N)\, 3^{-\delta' n}$, $\delta' \approx 0.33$. We then *sharpen this to the optimal rate*: refining the energy recursion by the pair-index residue modulo 3^m gives an exact linear transfer operator whose balanced top eigenvector forces equidistribution, yielding the unconditional level-one non-concentration $L(1,N)/M(N) = \frac{1}{3} + O(N^{-\gamma})$, $\gamma = \log_2(\sqrt{17}-3)$, and per-level $L(m,N)/M(N) \to 3^{-m}$ (concentration exponent $\delta = 1$) for every fixed m, with the exact finite spectrum of the mod-3 refinement (imbalance radius $\mu_2 = (5+\sqrt{17})/8$, level-independent). A closed-form ergodic-optimization lemma ($\max_\mu \int \log(2\sin^2 \pi x)\, d\mu = \log \frac{3}{2}$, unique maximizer the $\{\frac{1}{3}, \frac{2}{3}\}$ cycle) pins the resonance variationally. The one hypothesis left — that the imbalance radius equals μ_2 uniformly in m — is reduced by an exact twisted $3\times 3 \to 2\times 2$ cocycle reduction, the Galois-rational trace identity $D_m = \det Q_m = 8^{-P}$ ($\Phi_{3^m}(-1) = 1$), and a Schur–Cohn criterion to the single scalar inequality $T_m < 1$; this in turn reduces, by an exact Frobenius block-tiling bound and an exact Ramanujan equidistribution (whose units avoid $\{0, \frac{1}{3}, \frac{2}{3}\}$, arithmetically bypassing the $\omega = 1$ fixed-point obstruction that provably blocks every Lyapunov certificate), to the single integral inequality $\int_0^1 \varphi_{16} < 0$, verified to $> 20\sigma$ and covered exactly for $m \le 13$. Thus every arrow of the chain $[\,T_m < 1\,] \Rightarrow (\text{TM-}\ell^2, \delta = 1) \Rightarrow \text{K2} \Rightarrow (\text{MS}) \Rightarrow c_0 = \frac{1}{2}$ is proved except two named inputs: the pending interval certification of $\int \varphi_{16} < 0$ and the measured-favorable mass non-concentration (MULT-mass). The unconditional depth constant is unchanged at $c_0 \approx 0.124$; we do *not* claim $c_0 = \frac{1}{2}$ unconditionally. We state precisely what is proved, what is conjectured, and where the genuinely Collatz-hard quenched obstruction begins.

Additions in v6 (25 July 2026). Two developments postdate v5. (i) *Machine verification.* The three technical lemmas underlying the uniform $\delta = 1$ certificate — the Schur–Cohn stability step, Ramanujan exactness for the prime power 3^m, and the Jackson-kernel identity and weighted bound — have been formalised in Lean 4 with Mathlib. The development comprises 79 declarations and depends on no admitted statement: every theorem passes `#print axioms` resting only on `propext`, `Classical.choice` and `Quot.sound`. Since Mathlib contains no Fejér or Jackson kernel, the identity $\int_0^1 (\sin n\pi t/\sin \pi t)^4\, dt = n(2n^2+1)/3$ was built from character orthogonality upward; the proof reduces the integral to the cardinality $\#\{(i,j,k,l) \in [0,n)^4 : i+j = k+l\}$ and evaluates it by fibring. A kernel-only certified rational interval arithmetic is included for the finite residual. (ii) *An exact constant.* The carry-state transfer matrix appearing in recent work on Gowers-norm bounds for the Thue–Morse sequence is shown to coincide with the operator solved exactly in Section 8; its spectral radius is $(1+\sqrt{17})/8$, so the associated exponent is $\eta_0 = 1 - \log_2((1+\sqrt{17})/4) = 0.6429813631\ldots = 1-\beta$, the same algebraic number as the energy exponent of Theorem 7.4. See Section 10.

1. Introduction

The Collatz (3x+1) conjecture concerns the map $n \mapsto 3n+1$ (n odd), $n \mapsto n/2$ (n even). Its standard acceleration is the *Syracuse map* on odd integers,

$$\mathrm{Syr}(n) \;=\; \frac{3n+1}{2^{a(n)}}, \qquad a(n) = \nu_2(3n+1),$$

and we write $m_k(n) = \mathrm{Syr}^k(n)$. This paper studies the binary sum-of-digits function s_2 (equivalently the Hamming weight or popcount) along Syracuse iterates, through the exponential sums

$$S_k(\alpha; N) \;=\; \sum_{\substack{n \le N \\ n \text{ odd}}} \mathrm{e}\big(\alpha\, \mathrm{s}_2(m_k(n))\big), \qquad \mathrm{e}(x) = e^{2\pi i x}. \tag{1}$$

Equidistribution of $\mathrm{s}_2(m_k(n))$ in residue classes for *every* k along a *single* orbit (the "quenched" statement) is essentially the Syracuse binary-normality property, a 2-adic companion to Tao's 3-adic equidistribution conjecture for Syracuse random variables [20, 21]; it is open, and it is blocked by the non-arithmetic conjugacy obstruction of Bernstein–Lagarias [3]. The results below are *annealed* — averaged over the seed n — and are, to our knowledge, the first unconditional equidistribution statements for binary digits of Syracuse iterates.

1.1. Main results.

Theorem 1.1 (Fixed depth). *For every $\alpha \in \mathbb{R} \setminus \mathbb{Z}$ and every fixed $k \ge 1$,*

$$S_k(\alpha; N) = o_k(N), \qquad N \to \infty.$$

Consequently $\mathrm{s}_2(m_k(n))$ is equidistributed modulo q for every fixed $q \ge 2$, on average over odd seeds $n \le N$.

Theorem 1.2 (Growing depth, unconditional). *Let $\gamma(\alpha)$ be the Gelfond exponent (3) and $\theta_{\mathrm{prov}}(\alpha) = \gamma(\alpha)/(1-\gamma(\alpha))$. For every $\alpha \notin \mathbb{Z}$ and $\varepsilon > 0$, $S_k(\alpha; N) = o(N)$ uniformly for all*

$$k \;\le\; (c_0(\alpha) - \varepsilon) \log_2 N, \qquad c_0(\alpha) \;=\; \frac{1}{2 + \theta_{\mathrm{prov}}(\alpha) \log_2 3}.$$

The binding character is the parity character $\alpha = \frac{1}{2}$, for which $\gamma = \log_4 3$, $\theta_{\mathrm{prov}} = 3.82\ldots$, and $c_0 \approx 0.124$.

Three further constants calibrate Theorem 1.2. If the single-modulus digit-AP problem (Problem T1 below) were available with exponent $\theta = 1$, the same proof would give $c_0 = 1/(2 + \log_2 3) \approx 0.279$; the information-theoretic ceiling of the annealed statement is $c_0 = \frac{1}{2}$ (beyond depth $\frac{1}{2} \log_2 N$ the cylinder decomposition degenerates to single seeds and the statement becomes quenched); and numerically the cancellation in (1) persists to depth $\approx 2.41 \log_2 N = \log_2 N / \log_2(4/3)$, the orbit-absorption depth at which orbits reach the terminal cycle. Thus the rigorous constant is limited by proof technique, not by the truth: empirically the mean phase shows power-saving cancellation with onset length $O(1)$ in the modulus (Section 13).

1.2. The conditional path to the ceiling.

Our second contribution is a lossless structural reduction of the gap $0.124 \to \frac{1}{2}$. Define, for odd d and $\alpha \notin \mathbb{Z}$, the digit-AP sums and the transfer sum

$$S_a(L) = \sum_{n<L} \mathrm{e}\big(\alpha\, \mathrm{s}_2(a+dn)\big), \qquad T(\nu, d, \alpha) = \sum_{j \bmod d} \prod_{i<\nu} 4\cos^2\Big(\pi\Big(\alpha + \frac{2^i j}{d}\Big)\Big), \tag{2}$$

linked by the exact Parseval identity $\sum_{a \bmod d} |S_a|^2 = T(\nu, d, \alpha)/d$ for $N = 2^\nu = dL$ (Lemma 5.2). For $d = 3^m$ we prove that the transfer sum decomposes exactly along the $\times 2$-orbits of $\mathbb{Z}/3^m$ (one orbit per level, as 2 is a primitive root), that the full-period orbit products admit the closed cyclotomic evaluation of Theorem 6.1 — in particular $Q_m(\frac{1}{2}) = 9$ at *every* level m, so the Newman–Coquet resonance [13, 5] occurs exactly and only at level one — and that in the deep-window regime $r < m \log_2 3$ the window sums are *exactly* diagonal (Theorem 6.2). What remains is a single arithmetic inequality on signed binary (NAF) representations of multiples of 3^m (Conjecture 6.7); its truth, together with a mass non-concentration bound for Syracuse cylinder offsets (Conjecture 6.9, the repair of a count-multiplicity conjecture refuted in Remark 6.8), implies $c_0 = \frac{1}{2} - \varepsilon$ (Theorem 6.10). The measured shape of the kernel is recorded in Section 13: the excess is entirely resonance-shaped, $X_m(r) \approx 0.66 \cdot 2^{(2\gamma-1)r}/3^{m-1}$, giving $T = dN + \Theta(\log N)\, N^{2\gamma}$.

1.3. **What this does and does not say about Collatz.** Everything here is annealed. The quenched statement (fixed orbit, k up to the orbit length) is equivalent to a k-uniform seed$\times$depth bilinear bound which fails at the cylinder-aggregation step precisely because the word-to-offset map is non-arithmetic across words — the Bernstein–Lagarias obstruction [3], which we regard as the genuinely Collatz-hard boundary (Problem T2). No claim toward the Collatz conjecture is made. The interest of the annealed theorems is (i) they are, to our knowledge, the first unconditional digit-level equidistribution results for Syracuse iterates; (ii) they isolate, in Problems T1/T2 and Conjectures (MS)/(MULT-mass), exactly which parts of the problem are ordinary analytic number theory and which are Collatz-hard.

2. Preliminaries

2.1. **The Syracuse map in binary.** For odd n, $3n+1 = n + (n \ll 1) + 1$: an add-with-carry; the trailing bits of n determine $a(n) = \nu_2(3n+1)$, and $P(a(n) = j) = 2^{-j}$ for odd n drawn uniformly from a dyadic range (in the natural-density sense). Two exact facts anchor the digit viewpoint.

Lemma 2.1 (Weight invariance of the even steps)**.** $\mathrm{s}_2(\mathrm{Syr}(n)) = \mathrm{s}_2(3n+1)$.

Proof. Division by 2^a removes trailing zeros only; it shifts bit-length but never the multiset of ones. ☐

Lemma 2.2 (Piecewise-affine structure of m_k)**.** *Let $s = (a_1, \ldots, a_k) \in \mathbb{N}^k$, $\ell(s) = a_1 + \cdots + a_k$. On the 2-adic cylinder of odd seeds realizing valuation pattern s,*

$$m_k(n) \;=\; \frac{3^k n + c(s)}{2^{\ell(s)}}, \qquad c(s) \;=\; \sum_{i=1}^{k} 3^{\,k-i}\, 2^{\,a_1+\cdots+a_{i-1}}.$$

The cylinder is an arithmetic progression of seeds modulo $2^{\ell(s)+1}$, of density $2^{-\ell(s)-1}$ among odd residues; the total mass telescopes, $\sum_s 2^{-\ell(s)} = 1$ at each depth. Consequently the image $\{m_k(n)\}$ restricted to one cylinder is an arithmetic progression of odd difference 3^k.

Proof. Induction on k from $m_1(n) = (3n+1)/2^{a_1}$; composing affine maps with the stated valuations gives the formula, and the cylinder condition constrains n modulo $2^{\ell+1}$. ☐

Lemma 2.2 is the arithmetic handle on the iterate: the sum (1) splits into digit sums along APs of difference 3^k. The number of nonempty cylinders at depth k within $[1, N]$ is $\binom{\lfloor \log_2 N \rfloor}{k}(1 + o(1))$-many in the relevant range, which is $N^{H(c)(1+o(1))}$ for $k = c \log_2 N$, with H the binary entropy.

2.2. **The Gelfond exponent.** For $\alpha \notin \mathbb{Z}$ let

$$\gamma(\alpha) = \log_4\big(\text{spectral radius governing } \max_\beta \big|\sum_{m<M} \mathrm{e}(\alpha\,\mathrm{s}_2(m)+\beta m)\big|\big), \qquad \Big|\sum_{m<M} \mathrm{e}(\alpha\,\mathrm{s}_2(m)+\beta m)\Big| \le C_\alpha M^{\gamma(\alpha)}, \tag{3}$$

uniformly in β [8, 11]. For the parity character, $\gamma(\frac{1}{2}) = \log_4 3 = 0.79248\ldots$, sharp on the Newman–Coquet frequency $\beta = \pm\frac{1}{3}$ [13, 5]; $\gamma(\frac{1}{3}) = 0.7535\ldots$. Gelfond's theorem [8] gives, for each fixed odd d, equidistribution of s_2 along $n \equiv a \pmod d$ with a power saving whose constant depends on d; crucially, no known bound makes the exponent uniform in d (see Section 5).

3. Fixed depth: proof of Theorem 1.1

By Lemma 2.2, for fixed k split (1) over cylinders. Cylinders with $\ell(s) > A_k$ contribute at most $2^{-A_k}N$ in total mass; choose A_k large. The $O_k(1)$ remaining cylinders each contribute a sum $\sum \mathrm{e}(\alpha\ \mathrm{s}_2(\cdot))$ over an arithmetic progression of odd difference 3^k and length $\asymp_s N$. Gelfond's theorem [8] applies to each (in base 2 the coprimality condition $\gcd(d,2)=1$ is automatic for $d=3^k$), giving a power saving with a constant depending only on k. Summing the $O_k(1)$ cylinders and letting $A_k \to \infty$ slowly proves Theorem 1.1. □

Remark 3.1. The case $k=1$ is already the substantive novelty: equidistribution of $\mathrm{s}_2\big((3n+1)/2^{a(n)}\big)$ over odd $n \le N$. By Lemma 2.1 this equals $\mathrm{s}_2(3n+1)$: the theorem quantifies how the single shift-add-carry randomizes binary weight on average.

4. Growing depth: proof of Theorem 1.2

At depth $k = c\log_2 N$ the cylinder APs have typical length $N/2^{2k}$ (typical $\ell \approx 2k$) and difference $d = 3^k$; the proof of Theorem 1.1 survives as long as each AP is long enough for a d-uniform power saving. The engine is the finite Fourier expansion: with $G(M;\alpha,\beta) = \sum_{m<M} \mathrm{e}(\alpha\,\mathrm{s}_2(m)+\beta m)$,

$$\sum_{\substack{m<M\\ m\equiv a\ (d)}} \mathrm{e}(\alpha\,\mathrm{s}_2(m)) = \frac{1}{d}\sum_{j \bmod d} \mathrm{e}(-ja/d)\, G\big(M;\alpha,\tfrac{j}{d}\big), \qquad \Big|\cdot\Big| \le C_\alpha \frac{1}{d}\cdot d\cdot M^{\gamma} = C_\alpha M^{\gamma}. \tag{4}$$

The d-fold triangle inequality over the frequencies $\{j/d\}$ is the entire loss. A sum over an AP of length L inside $[0,dL)$ is nontrivial when $C(dL)^\gamma = o(L)$, i.e. $L \ge d^{\theta_{\text{prov}}+\varepsilon}$ with $\theta_{\text{prov}} = \gamma/(1-\gamma)$. Requiring $N/2^{2k} \ge (3^k)^{\theta_{\text{prov}}+\varepsilon}$ yields $k \le (c_0-\varepsilon')\log_2 N$ with $c_0 = 1/(2+\theta_{\text{prov}}\log_2 3)$, and the cylinder aggregation converges because the cylinder mass telescopes ($\sum_s 2^{-\ell(s)} = 1$) while the cylinder count $N^{H(c)}$ satisfies $H(c) < 1$ throughout $c < \frac{1}{2}$; a bounded-tail truncation in ℓ as in Theorem 1.1 completes the argument. At $\alpha = \frac{1}{2}$: $\theta_{\text{prov}} = 3.82$, $c_0 \approx 0.124$. □

Remark 4.1 (History of the constant)**.** An earlier version of this argument asserted $c_0 \approx 0.279$, corresponding to $\theta = 1$; that step assumed a digit-AP bound uniform in the offset which is *not* a theorem (the high-digit part of s_2 sweeps $\Theta(d)$ values along the AP rather than being constant over a low period). The corrected unconditional constant is $c_0 \approx 0.124$. Gelfond-type fixed-d savings [8, 10] have constants provably growing with d, and every $\theta < 1$ result in the literature — Fouvry–Mauduit [7], Müllner–Spiegelhofer [12], Spiegelhofer's level-of-distribution theorem [16] — is an average over the modulus d, the wrong quantifier for the single fixed $d = 3^k$ needed here.

5. The single-modulus problem T1 and the offset average

Open Problem 5.1 (T1)**.** Prove, for fixed odd $d \to \infty$ (specifically $d = 3^k$), a bound $\max_a \left|\sum_{n<L} \mathrm{e}(\alpha\, \mathrm{s}_2(a+dn))\right| = o(L)$ valid for $L \geq d^\theta$ with some $\theta < \theta_{\mathrm{prov}}(\alpha)$, uniformly in the offset a. Any $\theta < 3.82$ at $\alpha = \frac{1}{2}$ improves Theorem 1.2; $\theta = 1$ gives $c_0 = 0.279$; the numerics of Section 13 indicate the truth is $\theta = 0$ (onset length $O(1)$ in d).

Averaging over the offset removes half of the frequency loss but not enough:

Lemma 5.2 (Parseval identity)**.** *For $N = dL = 2^\nu$,* $\displaystyle\sum_{a \bmod d} |S_a(L)|^2 = \frac{1}{d} \sum_{j \bmod d} \left|G(N;\alpha,\tfrac{j}{d})\right|^2 = \frac{T(\nu,d,\alpha)}{d}$*, with T as in (2) and $|G_j|^2 = \prod_{i<\nu} 4\cos^2(\pi(\alpha + 2^i j/d))$.*

Proposition 5.3 (Offset-averaged bound)**.** $\left(\frac{1}{d}\sum_a |S_a|^2\right)^{1/2} \leq C\, d^{\gamma-1/2} L^\gamma$*, so the root-mean-square offset is nontrivial for $L \geq d^{\theta_{\mathrm{avg}}}$, $\theta_{\mathrm{avg}} = \frac{\gamma-1/2}{1-\gamma}$: $\theta_{\mathrm{avg}}(\frac{1}{2}) = 1.41$, $\theta_{\mathrm{avg}}(\frac{1}{3}) = 1.03$.*

Proof. Insert $\max_j |G_j| \leq CN^\gamma$ into Lemma 5.2. □

However, running the averaged bound through the cylinder aggregation of Section 4 by Cauchy–Schwarz costs a factor $\sqrt{M}$, $M = N^{H(c)}$ the cylinder count, and the resulting condition $H(c) < 2 - 2\gamma$ yields only $c \lesssim 0.086 < 0.124$: *the offset average alone does not improve the unconditional constant.* Its value is structural, via the exact analysis of the next section.

6. Exact structure of the transfer sum

Throughout, $d = 3^m$, $f(x) = 4\cos^2(\pi(\alpha + x))$, $N = 2^\nu$. Since 2 is a primitive root modulo 3^m, the nonzero frequencies $j/3^k$ of level m (i.e. $j = 3^{k-m}u$, $3 \nmid u$) form a *single* $\times 2$-orbit of length $P_m = 2 \cdot 3^{m-1}$.

Theorem 6.1 (Cyclotomic evaluation of full-period products)**.** *For odd q and any α, $\prod_{u \bmod q} 4\cos^2\left(\pi(\alpha + \frac{u}{q})\right) = 4\cos^2(\pi q\alpha)$. Hence the level-$m$ orbit product is*

$$Q_m(\alpha) = \frac{\cos^2(3^m \pi\alpha)}{\cos^2(3^{m-1}\pi\alpha)} \quad (\textit{when defined}), \qquad \textit{and} \qquad Q_m(\tfrac{1}{2}) = 9 \quad \textit{for every } m \geq 1.$$

Consequently the per-step growth rate $\lambda_m = \frac{\log_2 Q_m}{P_m}$ satisfies, at $\alpha = \frac{1}{2}$: $\lambda_1 = \log_2 3 = 2\gamma(\frac{1}{2})$ (the Newman–Coquet resonance, exactly and only at level one) and $\lambda_m \leq \frac{\log_2 9}{6} < 1$ for all $m \geq 2$; at $\alpha = \frac{1}{3}$: $Q_1 = 4$, $Q_m = 1$ for $m \geq 2$; and for generic α the products telescope, $\prod_{m \leq k} Q_m = \cos^2(3^k \pi\alpha)/\cos^2(\pi\alpha)$.

Proof. From $\prod_{u \bmod q}(1 + z\mathrm{e}(u/q)) = 1 - (-z)^q$ with $z = \mathrm{e}(\alpha)$ and $|1 + \mathrm{e}(x)|^2 = 4\cos^2(\pi x)$, using q odd. The primitive product is the ratio of the full products at $q = 3^m$ and $q = 3^{m-1}$. At $\alpha = \frac{1}{2}$ both sides vanish (the $u \equiv 0$ factor); instead use $\prod_{u=1}^{q-1} 2\sin(\pi u/q) = q$, whence $\prod_{3 \nmid u} 4\sin^2(\pi u/3^m) = (3^m)^2/(3^{m-1})^2 = 9$, and $f_{1/2}(x) = 4\sin^2(\pi x)$. □

Theorem 6.2 (Exact deep-window diagonal)**.** *For every α and every window length $r < m\log_2 3$,*

$$\sum_{u \bmod 3^m} \prod_{i<r} f\left(\frac{2^i u}{3^m}\right) = 3^m\, 2^r \qquad \textit{exactly.}$$

Proof. Expand $\prod_i \big(2 + \mathrm{e}(\alpha + 2^i u/3^m) + \mathrm{e}(-\alpha - 2^i u/3^m)\big)$ into signed-digit configurations $\varepsilon \in \{0, \pm 1\}^r$ with phases $\mathrm{e}(u\, W(\varepsilon)/3^m)$, $W(\varepsilon) = \sum \varepsilon_i 2^i$. Summation over all u annihilates every configuration with $3^m \nmid W$. Distinct signed powers of two never sum to zero, and $0 < |W| < 2^r < 3^m$ excludes divisibility, so only $\varepsilon = 0$ survives, with weight 2^r per residue. □

In the normalization $g(x) = 2\sin^2(\pi x)$ (so that at $\alpha = \frac{1}{2}$ each factor is $f/2$), the primitive level sums $B_t(r) = \sum_{v \text{ prim } (3^t)} \prod_{i<r} g(2^i v/3^t)$ admit three further exact statements.

Theorem 6.3 (Exact resonant term and primitive decomposition)**.** *(i) $g \equiv \frac{3}{2}$ identically on the level-one orbit $\{\frac{1}{3}, \frac{2}{3}\}$, hence $B_1(r) = 2 \cdot (\frac{3}{2})^r$ exactly for every $r \geq 1$; in particular the kernel constant below is exactly $\frac{2}{3}$. (ii) The $(\frac{3}{2})^r$ term cancels from every primitive difference: $A_m(r) = \sum_{u \bmod 3^m} \prod_{i<r} g = 2(\frac{3}{2})^r + \sum_{t\geq 2} B_t(r)$. (iii) For $r \geq P_t = 2 \cdot 3^{t-1}$, $t \geq 2$: $B_t(r) \leq 9C\, P_t(\frac{3}{4})^{P_t}$ (exponentially suppressed; $O(1)$ in total), while for $r < t\log_2 3$, $B_t(r) = 2 \cdot 3^{t-1}$ exactly. Moreover $B_t(r) \geq 0$ always, since $g \geq 0$.*

Consequently the full conjectural bound reduces to a *one-sided, per-level* estimate:

Conjecture 6.4 (K2)**.** *$B_t(r) \leq C \cdot 3^t$ uniformly in the mid-window $t \log_2 3 \leq r < 2 \cdot 3^{t-1}$.*

Conjecture 6.4 is equivalent to Conjecture 6.7 below (with the exact constant $\frac{2}{3}$), and no sign cancellation is required anywhere: only an upper bound on a positive quantity. Measured: $0 \leq B_t/(2 \cdot 3^{t-1}) \leq 1.43$ throughout, with the constant apparently tending to 1; the sharper extended measurement below gives $B_t \leq 1.04 \cdot 3^t$ for $t \leq 14$.

The primitive sums B_t are governed by a finite transfer operator whose spectrum we can compute exactly. On $\mathbb{C}[\mathbb{Z}/3^t]$ let $(M_t h)(u) = g(u/3^t)\, h(2u \bmod 3^t)$, so that $B_t(r) = \mathbf{1}^\top \mathrm{block}_t^r \mathbf{1}$ where block_t is the restriction of M_t to the units.

Theorem 6.5 (Exact finite spectrum of M_t)**.** *Write $P_s = 2 \cdot 3^{s-1}$ and $\mu_P = \{\zeta : \zeta^P = 1\}$.*

(i) *The valuation sectors $V_k = \mathrm{span}\{e_u : \nu_3(u) = k\}$ are each M_t-invariant.*

(ii) *On the units V_0, the map $u \mapsto 2u$ is a* single *cycle of length P_t (as 2 is a primitive root modulo 3^t), so* block_t *is the cyclic weighted shift with weights $w_j = g(2^j u_0/3^t)$ along the orbit, and $B_t(r) = \sum_{j=0}^{P_t-1} \prod_{i=0}^{r-1} w_{j+i}$ (indices mod P_t).*

(iii) *The spectrum of* block_t *is the full circle of radius*

$$r_t = \Big(\prod_j w_j\Big)^{1/P_t} = \Big(\frac{9}{2^{P_t}}\Big)^{1/P_t} = \frac{9^{1/P_t}}{2},$$

strictly decreasing with $r_t \downarrow \frac{1}{2}$ and $r_t \leq 9^{1/6}/2 = 0.72113 < 1$ for all $t \geq 2$. Consequently

$$\mathrm{spec}(M_t) = \{\pm\tfrac{3}{2}\} \;\cup\; \bigcup_{s=1}^{t} r_s \cdot \mu_{P_s} \;\cup\; \{0\},$$

the top pair $\pm\frac{3}{2}$ carried by the level-one suborbit $\{3^{t-1}, 2 \cdot 3^{t-1}\}$; and for each fixed t, $B_t(r) \to 0$ geometrically as $r \to \infty$.

Proof. (i) Since 2 is a unit modulo 3^t, $\nu_3(2u \bmod 3^t) = \nu_3(u)$, and multiplication by $g(u/3^t)$ does not move support; hence each V_k is invariant. (ii) 2 generates $(\mathbb{Z}/3^t)^\times$, of order $\varphi(3^t) = 2 \cdot 3^{t-1} = P_t$, so the units form one $\langle 2 \rangle$-orbit; on it M_t is the weighted cyclic shift, and $\mathbf{1}^\top \mathrm{block}_t^r \mathbf{1}$ is the sum of its length-r window products. (iii) A length-P weighted cyclic shift has characteristic polynomial $\lambda^P - \prod_j w_j$, so its eigenvalues are the P-th roots of $\prod_j w_j$, all of modulus $(\prod_j w_j)^{1/P}$. Here $\prod_j w_j = \prod_{3\nmid v} g(v/3^t) = \prod_{3\nmid v} \frac{1}{2} \cdot 4\sin^2(\pi v/3^t) = 2^{-P_t} \cdot 9$ by the cyclotomic evaluation of Theorem 6.1 ($\prod_{3\nmid v} 4\sin^2(\pi v/3^t) = (3^t)^2/(3^{t-1})^2 = 9$), giving $r_t = 9^{1/P_t}/2$. The valuation-k sector reproduces the level-$(t-k)$ primitive block and $u=0$ contributes 0 (as $g(0)=0$), so $\mathrm{spec}(M_t)$ decomposes over levels as stated; level one has $P_1 = 2$, $r_1 = 9^{1/2}/2 = \frac{3}{2}$, roots $\{\pm 1\}$, i.e. the pair $\pm\frac{3}{2}$. Finally $r_t < 1$ gives $B_t(r) \to 0$ by Gelfand. □

Theorem 6.5 settles the qualitative danger flagged by K2 — an eigenvalue creeping toward 1: the spectral radius is uniformly ≤ 0.7211 with a constant gap below 3/2. What it does *not* settle is the transient: block_t is extremely non-normal, so $r_t < 1$ does not bound $\sup_r B_t(r)/P_t$; that supremum is a maximal sub-arc product of the cyclic weights, measured ≤ 1.56 (hence $B_t \le 1.04 \cdot 3^t$) for $t \le 14$ with no growth trend, but not proved uniform. This transient is precisely the mid-window kernel of K2.

Remark 6.6 (Cycle census: correction to the resonance picture)**.** The resonance is the doubling 2-cycle $\{\frac{1}{3}, \frac{2}{3}\}$, on which $g \equiv \frac{3}{2}$. An exhaustive census of all doubling cycles of period ≤ 20 (111,012 cycles) corrects a natural but false premise: $\{\frac{1}{3}, \frac{2}{3}\}$ is *not* the unique cycle whose per-period geometric mean $\mathrm{gm}(C) = \left(\prod_{x\in C} g(x)\right)^{1/|C|}$ exceeds 1 (2,783 such cycles exist, e.g. $\mathrm{gm} \approx 1.30$ on a period-7 cycle). What is true — and what the argument actually needs — is that $\{\frac{1}{3}, \frac{2}{3}\}$ is the unique *maximizer*: $\max_C \mathrm{gm}(C) = \frac{3}{2}$ exactly, uniquely attained there (runner-up ≈ 1.41). Since $\frac{3}{2} < 2$, each cycle's geometric window-sum $\sum_j (\mathrm{gm}/2)^j$ converges, so no single resonance breaks the 3^t budget; the obstruction is the sum over the infinitely many sub-dominant cycles — an ergodic-optimization quantity, $\sup_\mu \exp\int \log g\, d\mu = \frac{3}{2}$ over doubling-invariant μ.

Two further structural facts inform Conjecture 6.4. Via Ramanujan sums, $B_t(r) = \sum_\varepsilon (-\frac{1}{2})^{s(\varepsilon)} c_{3^t}(W(\varepsilon))$, a signed, weight-discounted count over signed binary representations of multiples of 3^{t-1}. The *unsigned* weighted representation count $\sum_{\text{reps of } W} 2^{-s}$ tends to 1 for every odd W (padding tails), so minimal-weight lower bounds alone cannot prove K2 — the alternation $(-1)^s$ is essential. In the infinite-window limit the signed count $\rho(W) = \sum_{\text{reps}} (-\frac{1}{2})^s$ satisfies $\rho(0) = 1$, $\rho(2W) = \rho(W)$, $\rho(W) = -\frac{1}{2}[\rho(\frac{W-1}{2}) + \rho(\frac{W+1}{2})]$ for odd W, and is precisely the W-th Fourier coefficient of the Thue–Morse spectral measure, the non-dissociate Riesz product $\prod_i (1 - \cos(2\pi 2^i x))$ [14, 1]. We correct here a claim of the previous version: the pointwise bound $|\rho(W)| \le 2^{-s_{\mathrm{NAF}}(W)}$ is *false* — already $\rho(3) = \frac{1}{3} > 2^{-2}$, and the purported induction breaks precisely when the two branch depths of the recursion coincide. What is true is only an empirical, unproved Lyapunov rate of the averaging recursion: numerically $-\log_2 |\rho(W)| \approx 0.73\, s_{\mathrm{NAF}}(W)$ (measured to $W \le 10^7$, validated against exact rational arithmetic). Conjecture 6.4 at infinite window is thus a boundedness statement for Ramanujan-weighted sums of Thue–Morse Fourier coefficients along the multiples of 3^{t-1}; the residual arithmetic input is a *slab* bound

$$\sum_{1\le |w| \le X} 2^{-s_{\mathrm{NAF}}(w\cdot 3^{t-1})} \le C \qquad (X \le 2^{\kappa_0 t}), \tag{5}$$

on signed binary weights of multiples of powers of 3. The *unrestricted* sum is not bounded: a random-digit model, validated against exact computation over six orders of magnitude, gives $\sum_{|w|\le X} 2^{-s_{\mathrm{NAF}}(w\cdot 3^{t-1})} \asymp X^{1+\log_2(5/6)} = X^{0.737}$, whence boundedness can hold only on the slab $X \le 2^{\kappa_0 t}$ with sharp (model-)constant

$$\kappa_0 \;=\; \frac{-\log_2(5/6)\,\log_2 3}{1 + \log_2(5/6)} \;=\; 0.566.$$

The relevant growth of the underlying weight is $s_{\mathrm{NAF}}(3^n)/n \to 0.527$, measured to $n = 10^4$; its linearity is unproved — the only proved lower bounds are logarithmic-type and concern unsigned digits [18, 15, 6], and the joint base-2/base-3 digit machinery closest to the signed question is that of Spiegelhofer [17].

Writing $\nu = qP_m + r$, the level-m contribution to T equals Q_m^q times the primitive window sum at remainder r; Theorems 6.1, 6.2 and 6.3 make every piece exact except the *mid-window regime* $m \log_2 3 \le r < P_m$, where the window sum exceeds the diagonal by

$$X_m(r) \;=\; \sum_{\substack{\varepsilon \in \{0,\pm 1\}^r \setminus \{0\} \\ 3^m \mid W(\varepsilon)}} (-1)^{s(\varepsilon)}\, 2^{-s(\varepsilon)} \qquad (\alpha = \tfrac{1}{2}), \tag{6}$$

$s(\varepsilon)$ the number of nonzero digits: a purely arithmetic, signed, weight-discounted count of signed-binary representations of multiples of 3^m. The weight-2 terms require $3^{m-1} \mid (i - i')$ and total $O(r^2/3^{m-1})$; the general term lies in the Senge–Straus/Stewart circle of base-2/base-3 digit independence [15, 18].

Conjecture 6.7 (MS: mean-square / kernel bound)**.** *For $\alpha = \frac{1}{2}$ and $m \log_2 3 \le r < P_m$,*

$$X_m(r) \;\le\; C\,\frac{2^{(2\gamma-1)r}}{3^{m-1}}, \qquad \textit{equivalently} \qquad T(\nu, 3^k, \alpha) \;\le\; C\left(dN + N^{2\gamma}\right)\log N.$$

For nonresonant α (e.g. $\alpha = \frac{1}{3}$) the $N^{2\gamma}$ term is absent.

By Theorem 6.3 the leading term of the kernel is no longer conjectural: the constant is exactly $\frac{2}{3}$ (i.e. $X_m(r) = 2(\frac{3}{2})^r/3^m + [\sum_{t\ge 2} B_t - (3^m - 3)]/3^m$), and Conjecture 6.7 is equivalent to the one-sided Conjecture 6.4. The measured residual matches: exponents $2^{2\gamma-1}$ (per r) and 3^{-1} (per m) exact to three digits, excess entirely resonance-shaped. A sign oscillation occurs just past the deep-window threshold ($X < 0$), consistent with the measured transient of the primitive sums B_t.

The offsets are provably *not* equidistributed: $c(s) \equiv 2^{a_1+\cdots+a_{k-1}} \pmod 3$, so $3 \nmid a(s)$ always and the offsets occupy exactly two-thirds of the residues (observed saturating exactly at $2d/3$). The relevant hypothesis is a non-concentration of the *mass*, not of a count.

Remark 6.8 (Refutation of the count multiplicity, and its repair)**.** An earlier version of this paper conjectured that within each dyadic class of lengths $\ell(s)$ the *count* multiplicity $m(a) = \#\{s : a(s) = a\}$ is polylogarithmic. This is *false*. The offsets $a(s)$ lie in the fixed unit group $(\mathbb{Z}/3^k)^\times$, of size $D = 2 \cdot 3^{k-1}$, whereas a length class holds $\Theta(\ell^{k-1})$ words; pigeonhole forces $m(a) = \Theta(\ell^{k-1})$, a power of the class size, not polylog. (Exact enumeration confirms saturation of D by $\ell \approx 24$–31, with multiplicity jumps of $\times 7$–$\times 180$ across a single dyadic step.) The repair is to weight cylinders by the mass $2^{-\ell(s)}$ they actually carry in the aggregation — a cylinder contributes $S_{a(s)}(L_s)$ with $|S_{a(s)}(L_s)| \le L_s = N 2^{-\ell(s)-1}$, so its weight is exactly its mass. Set

$$W_j(a) = \sum_{\substack{s:\ \ell(s) \in [2^j, 2^{j+1}) \\ a(s) = a}} 2^{-\ell(s)}, \qquad M_j = \sum_a W_j(a), \qquad \kappa_j = \Big(\frac{\sum_a W_j(a)^2}{M_j^2}\Big)\, D \ \ (\ge 1).$$

κ_j is the mass-collision factor relative to a flat spread over the units ($\kappa_j = 1$ iff the mass is perfectly flat; D/κ_j is the effective residue support). It is a *second-moment* functional: the Cauchy–Schwarz step contracts through $\sum_a W_j(a)^2$, not through $\max_a W_j(a)$. The distinction is essential — $\max_a W_j(a)$ is exponentially large (the single all-ones cylinder $s = (1, \ldots, 1)$ has offset $3^k - 1$ and mass 2^{-k}, exceeding the flat share M_j/D by $\approx (3/2)^k$), yet one heavy residue adds only $\approx M_j^2/D$ to $\sum_a W_j(a)^2$, i.e. $O(1)$ to κ_j. Measured (exact enumeration, $k=8,10,12$): on the mass-dominant class $\kappa_j = 1.4, 2.2, 4.8$ (small, and *decreasing* as coverage of the class improves); on the tiny under-saturated shortest class κ_j is larger and grows with k, but that class carries $o(1)$ of the mass and its few long-AP cylinders are handled directly by the fixed-depth bound.

Conjecture 6.9 (MULT-mass: offset mass non-concentration)**.** *For $k = c\log_2 N$ with $c < \frac{1}{2}$, the mass-collision factor is polylogarithmic uniformly in the dyadic length class:*

$$\kappa_j = \Big(\frac{\sum_a W_j(a)^2}{M_j^2}\Big)\cdot 2\cdot 3^{k-1} \;\le\; P(k,\log N) \qquad (P \text{ a fixed polynomial}),$$

equivalently the effective residue support satisfies $M_j^2/\sum_a W_j(a)^2 \;\ge\; 2\cdot 3^{k-1}/P$.

Theorem 6.10 (Conditional depth at the ceiling)**.** *Assume Conjectures 6.7 and 6.9. Then for every $\alpha \notin \mathbb{Z}$ and $\varepsilon > 0$, $S_k(\alpha; N) = o(N)$ uniformly for $k \le (\frac{1}{2}-\varepsilon)\log_2 N$.*

Proof sketch. Write $S_k = \sum_\ell \sum_{a \bmod d} m_\ell(a) S_a(L_\ell)$, $L_\ell = N2^{-\ell-1}$, and apply Cauchy–Schwarz at each fixed length against $\sum_a |S_a(L_\ell)|^2 = T/d$ (Lemma 5.2), renormalising the count collision as $\sum_a m_\ell(a)^2 = (C_\ell^2/D)\kappa_\ell$, $C_\ell = \binom{\ell-1}{k-1}$. Inserting (MS) and summing over ℓ, the factor $C_\ell\sqrt{L_\ell}$ carries $2^{-\ell/2}$, so the mass reappears: the diagonal sum is $\lesssim \sqrt{\bar\kappa}\,\sqrt{N}\,(2.414/\sqrt{3})^k$ and the resonant sum $\lesssim \sqrt{\bar\kappa}\,N^\gamma(2^{-\gamma}/(1-2^{-\gamma}))^k 3^{k(\gamma-1/2)}$, both $o(N)$ for every $c < \frac{1}{2}$ once $\bar\kappa = \max_j \kappa_j \le P$ is polylog (MULT-mass). The binding constraint is therefore the information-theoretic cylinder-nonemptiness wall $\ell \le \log_2 N$, i.e. $H(c) < 1 \Leftrightarrow c < \frac{1}{2}$; the polylog κ and the (MS) logarithms are absorbed into $N^{o(1)}$, with the tightest pinch near $c \approx \frac{1}{4}$ (resonant regime, where $H(c) < 2 - 2\gamma + c\log_2 3$). $\square$

7. The Thue–Morse route: non-concentration on 3-adic sublattices

The slab bound (5) — and with it the entire conditional path to $c_0 = \frac{1}{2}$ — is governed by the ℓ^2 mass of the Thue–Morse coefficients $\rho(W)$ on the sublattices $3^n\mathbb{Z}$. This is the most promising remaining analytic target because it is a question about a *fixed classical measure*, with an explicit functional equation, in which no Collatz object appears. Write

$$M(N) = \sum_{1\le |W|\le N} \rho(W)^2, \qquad L(n,N) = \sum_{\substack{1\le |W|\le N\\ 3^n | W}} \rho(W)^2.$$

Here $\sum_W \rho(W)^2\,\mathrm{e}(W\theta)$ is the Fourier series of the autocorrelation $\gamma = \mu * \tilde\mu$ of the Thue–Morse measure μ (Wiener), and $L(n,\cdot)$ samples γ along the 3-adic rationals $\{j/3^n\}$, which are *generic* (non-dyadic) points of γ because the self-similarity of μ is dyadic.

That the kernel B_t of Section 6 is *literally* a Thue–Morse lattice mass is an exact identity. Writing $z = \mathrm{e}(v/3^t)$ and $Q_r(z) = \prod_{i<r}(1-z^{2^i}) = \sum_{m<2^r} \varepsilon_m z^m$, where $\varepsilon_m = (-1)^{s_2(m)}$ is the Thue–Morse sequence, and $c_k(t,r) = \sum_{m<2^r,\, m\equiv k\,(3^t)} \varepsilon_m$, $S_t(r) = \sum_{k \bmod 3^t} c_k(t,r)^2$:

Theorem 7.1 (Kernel–Thue–Morse identity)**.** *For all $t \ge 1$, $r \ge 0$,*

$$B_t(r) \;=\; 2^{-r}\big[\,3^t\,S_t(r) - 3^{t-1}\,S_{t-1}(r)\,\big].$$

Proof. Since $|1-\mathrm{e}(x)|^2 = 2g(x)$ with $g(x) = 2\sin^2(\pi x)$, one has $\prod_{i<r} g(2^i v/3^t) = 2^{-r}|Q_r(z)|^2$. A residue v is a unit iff z is a primitive 3^t-th root of unity; summing $|Q_r|^2$ by Parseval on $\mathbb{Z}/3^t$ gives $\sum_{v \bmod 3^t}\prod_{i<r} g = 3^t S_t(r)$, and subtracting the level-$(t-1)$ (imprimitive) part $3^{t-1}S_{t-1}(r)$ isolates the primitive sum $B_t(r)$. $\square$

Thus the residual arithmetic input is a pure statement about the classical Thue–Morse measure: the dyadic-self-similar spectral measure must not concentrate on the triadic lattice $\{j/3^t\}$, after the consecutive-level subtraction. The autocorrelation γ (hence μ) is moreover invariant under doubling, which is what makes the whole $\times 2 \times 3$ circle of ideas available.

Theorem 7.2 (Doubling invariance)**.** *Let $T : x \mapsto 2x \bmod 1$. Then $T_*\mu = \mu$ and $T_*\sigma = \sigma$, where $\sigma = \mu * \tilde{\mu}$ (so $\widehat{\sigma}(k) = \rho(k)^2 \geq 0$).*

Proof. For any finite measure ν on $\mathbb{T}$, $\widehat{T_*\nu}(k) = \int e(-k \cdot 2x)\, d\nu = \widehat{\nu}(2k)$; Fourier coefficients determine ν, so $T_*\nu = \nu \iff \widehat{\nu}(2k) = \widehat{\nu}(k)$ for all k. Here $\widehat{\mu}(k) = \rho(k)$ with $\rho(2k) = \rho(k)$ (definitional), giving $T_*\mu = \mu$; and $\widehat{\sigma}(k) = \rho(k)^2$ with $\rho(2k)^2 = \rho(k)^2$, giving $T_*\sigma = \sigma$. □

Open Problem 7.3 (TM-ℓ^2: non-concentration)**.** Prove $L(n, N) \leq C\, M(N)\, 3^{-n\delta}$ for some $\delta > 0$, uniformly in N.

We report the following measurements (exact recursion, float64 validated against rational arithmetic to 1.4×10^{-17}, single $O(N)$ pass to $N = 10^7$):

(i) *ℓ^2 growth.* $M(N) \sim N^\beta$ with $\beta \approx 0.36$, i.e. a correlation dimension $D_2 = 1 - \beta \approx 0.64$ for μ (consistent with its known singular-continuous spectral type and the scaling analysis of Baake–Gähler–Grimm [2]). *Now proved:* $M(N) \geq \frac{2}{9}N^\beta$ with $\beta = \log_2 \frac{1+\sqrt{17}}{4} = 0.357019\ldots$ exactly (Theorem 7.4), matching the measured slope; the matching upper bound gives $M(N) = \Theta(N^\beta)$.

(ii) *Non-concentration, uniform.* The excess mass factor $F(n) = \big(L(n, N)/M(N)\big)/3^{-n}$ stays in $[0.40, 1.13]$ and decreases with n, yielding the uniform *measured* bound
$$L(n, N) \;\leq\; 1.13\, M(N)\, 3^{-n} \qquad (\text{all } n,\ N \leq 10^7),$$
so Problem 7.3 holds numerically with $\delta = 1$ and a small constant: the mass demanded by $3^n \mid W$ is comparable to its uniform density share 3^{-n}, with only a bounded, geometrically shrinking excess on the 0-class.

(iii) *Per-level contraction.* $L(n, N) \approx \lambda\, L(n-1, N)$ with $\lambda \approx 0.30 < \frac{1}{3}$, stable across levels and across N (giving $\delta = -\log_3 \lambda \approx 1.08$). *Now proved (unconditionally):* $L(n, N) \leq C\, M(N)\, 3^{-\delta' n}$ with $\delta' \approx 0.33 > 0$, i.e. $\lambda \leq 3^{-0.33} \approx 0.70 < 1$ (Theorem 7.5), short of the measured sharp rate.

Two of these three measured facts are now *proved*. The growth exponent (i) is exactly a transfer-operator eigenvalue, and admits an unconditional power lower bound.

Theorem 7.4 (Energy lower bound)**.** *Let $\beta = \log_2 \frac{1+\sqrt{17}}{4} = 0.3570186\ldots$. For every integer $N \geq 1$,*
$$M(N) \;>\; \frac{7 + \sqrt{17}}{18}\, N^\beta - 1 \;\geq\; \frac{2}{9}\, N^\beta.$$

Proof. Read the recursion on consecutive pairs $p_k = (\rho(k), \rho(k+1))$. With $a = \rho(k)$, $b = \rho(k+1)$, $y = -\frac{1}{2}(a+b)$ one has $\rho(2k) = a$, $\rho(2k+1) = y$, $\rho(2k+2) = b$, so $p_{2k} = (a, y)$, $p_{2k+1} = (y, b)$, and as k ranges over $0, \ldots, 2^n - 1$ the children range bijectively over $0, \ldots, 2^{n+1} - 1$. Fix
$$\alpha = \frac{5 + \sqrt{17}}{4}, \qquad \phi(a, b) = \alpha(a^2 + b^2) + ab, \qquad \lambda = \alpha - 1 = \frac{1 + \sqrt{17}}{4}.$$
Eigen-identity: $\phi(a, y) + \phi(y, b) = \lambda\, \phi(a, b)$ for all a, b. Indeed the left side is $\alpha(a^2 + b^2) + 2\alpha y^2 + y(a + b)$; since $a + b = -2y$ this is $\alpha(a^2 + b^2) + (2\alpha - 2)y^2$, and $y^2 = \frac{1}{4}(a^2 + b^2) + \frac{1}{2}ab$ gives $\frac{3\alpha - 1}{2}(a^2 + b^2) + (\alpha - 1)ab$, which equals $\lambda\phi(a, b)$ iff $\frac{3\alpha-1}{2} = \lambda\alpha$ and $\alpha - 1 = \lambda$, i.e. iff

$2\alpha^2-5\alpha+1=0$; the root $\alpha=\frac{5+\sqrt{17}}{4}$ yields $\lambda=\frac{1+\sqrt{17}}{4}$. *Block identity:* summing over k and using the child bijection, $\Phi_n:=\sum_{k=0}^{2^n-1}\phi(\rho(k),\rho(k+1))=\lambda^n\Phi_0=\lambda^n\phi(1,-\frac{1}{3})=\lambda^n\frac{19+5\sqrt{17}}{18}$. *Comparison:* $(\alpha+\frac{1}{2})(a^2+b^2)-\phi(a,b)=\frac{1}{2}(a-b)^2\geq 0$, so with $S(N)=\sum_{W=1}^{N}\rho(W)^2=M(N)/2$,

$$\Phi_n\leq(\alpha+\tfrac{1}{2})\sum_{k=0}^{2^n-1}(\rho(k)^2+\rho(k+1)^2)\leq(\alpha+\tfrac{1}{2})\big(1+2S(2^n)\big),$$

whence $1+2S(2^n)\geq\Phi_n/(\alpha+\frac{1}{2})=\frac{3+\sqrt{17}}{9}\lambda^n$, i.e. $S(2^n)\geq C\lambda^n-\frac{1}{2}$ with $C=\frac{3+\sqrt{17}}{18}$. For general N, take $n=\lfloor\log_2 N\rfloor$, so $2^n>N/2$ and $\lambda^n=(2^n)^\beta>\lambda^{-1}N^\beta$; monotonicity of S gives $M(N)=2S(N)>\frac{2C}{\lambda}N^\beta-1=\frac{7+\sqrt{17}}{18}N^\beta-1$. Since $\frac{7+\sqrt{17}}{18}-\frac{2}{9}=\frac{3+\sqrt{17}}{18}$, the clean form $M(N)\geq\frac{2}{9}N^\beta$ follows for $N\geq 14$, and holds for $N\leq 13$ by direct computation (the minimum of $M(N)/N^\beta$ is $\frac{2}{9}$, attained at $N=1$). □

The exponent β is the top eigenvalue of the pair substitution on the space of symmetric quadratic forms $\{a^2+b^2,ab\}$; it matches the measured 0.36 exactly, and the same form gives a matching upper bound $S(2^n)\leq\frac{\phi(1,-1/3)}{\alpha-1/2}\lambda^n$, so $M(N)=\Theta(N^\beta)$. Feeding Theorem 7.4 into the renormalization of measurement (iii) removes its one hypothesis:

Theorem 7.5 (Unconditional per-level contraction)**.** *There are $C<\infty$ and $\delta'\approx 0.33>0$ with $L(n,N)\leq C\,M(N)\,3^{-\delta' n}$ for all n,N; equivalently a geometric per-level contraction $\lambda\leq 3^{-\delta'}\approx 0.70<1$.*

Proof sketch. Convexity of squares on the odd recursion gives, for the class masses $m^{\leq N}(r)=\sum_{|W|\leq N,\,W\equiv r\,(3^n)}\rho(W)^2$, the componentwise transfer inequality $m^{\leq N}\leq T\,m^{\leq N/2}$, where $(Tm)(r)=m(2^{-1}r)+\frac{1}{2}m(2^{-1}(r-1))+\frac{1}{2}m(2^{-1}(r+1))$ (halving permutes residues since 2 is a unit mod 3^n). The dual chain $s\mapsto 2s+\varepsilon$ with $\varepsilon\in\{0,\pm1\}$ of probabilities $\{\frac{1}{2},\frac{1}{4},\frac{1}{4}\}$ has the uniform law as its *exact* stationary distribution (so the 0-class mass tends to 3^{-n}) and mixes in $\Theta(n)$ steps. Iterating K times, $L(n,N)=m^{\leq N}(0)\leq 2^K M(N/2^K)(3^{-n}+D_K)$; with the now-proved growth $M(N/2^K)\leq c\,M(N)2^{-\beta K}$ (Theorem 7.4, $\beta>0$) the prefactor becomes $2^{K(1-\beta)}$, and optimizing $K\asymp n$ yields $L(n,N)\leq C\,M(N)\,3^{-\delta' n}$ with $\delta'=1-c(1-\beta)\log 2/\log 3\approx 0.33$. □

This is short of the *measured* $\lambda\approx 0.30$ ($\delta\approx 1$): the convexity step loses a fixed factor 2 per scale while the chain needs $\Theta(n)$ scales to mix, structurally capping this route near $\delta'\approx 0.33$–0.39. It proves non-concentration (Problem 7.3 with $\delta=\delta'>0$) but not the sharp rate needed for K2. The remaining sharp input is the uniform *transient* bound of Theorem 6.5 — equivalently, that the maximal sub-arc products of the cyclic weights w_j (the hot windows of g along the $\times 2$ orbit modulo 3^t) are bounded after normalization. By Theorem 7.1 this is exactly (TM-ℓ^2); four independent attacks (signed NAF/Ramanujan sums, the Thue–Morse lattice masses, the finite operator spectra of Theorem 6.5, and the hot-window decomposition) now provably reduce to this single non-concentration fact for the dyadic Thue–Morse measure at triadic rationals.

Assuming (TM-ℓ^2), Cauchy–Schwarz against the sublattice mass yields the slab (5), and the program closes along the chain

$$\text{(TM-}\ell^2\text{)}\implies\text{(SW)-slab}\xRightarrow{+\text{tail sign-canc. (open)}}\text{K2}\implies(\text{MS},\ \tfrac{2}{3}\text{ exact})\xRightarrow{+\text{(MULT-mass)}}c_0=\tfrac{1}{2}-\varepsilon.$$

Every arrow past the first is established or one-sided-explicit; in the v4 form the two genuinely open links were (TM-ℓ^2) itself and the sign-cancellation of the mid-window tail feeding K2.

Both are resolved on the sharp route developed in Sections 8–9: the triadic eigen-form proves (TM-ℓ^2) at the optimal strength $\delta = 1$ (unconditionally at level one, per level for every fixed m), bypassing the mid-window sign-cancellation entirely, and reduces the single remaining uniform-in-m hypothesis to one certifiable scalar inequality. We still emphasize that this is a route, not a finished unconditional theorem: the unconditional depth constant remains $c_0 \approx 0.124$ (Theorem 1.2), and the sharp $c_0 = \frac{1}{2}$ path is conditional on the two named inputs recorded in Section 9.4 — but the per-level contraction (Theorem 7.5) and growth (Theorem 7.4) are unconditional, the spectral radius is uniformly below 1 (Theorem 6.5), and the non-concentration itself is now proved sharp per level.

8. The triadic eigen-form: sharp per-level non-concentration

The per-level contraction of Theorem 7.5 was proved only at the sub-optimal rate $\delta' \approx 0.33$, short of the measured $\delta = 1$ (measurement (iii) of Section 7). We now close that gap. Refining the energy recursion of Theorem 7.4 by the residue of the pair index modulo 3^m produces an *exact linear transfer operator* whose top eigenvector is perfectly balanced across residues, forcing equidistribution of the sublattice mass. Level one is unconditional; every fixed level is proved; the sole uniform-in-m hypothesis is reduced, in Section 9, to a single scalar inequality.

8.1. **The refined transfer operator.** Recall (Theorem 7.4) the consecutive pair $p_k = (\rho(k), \rho(k+1))$ and the substitution $p_{2k} = (a, y)$, $p_{2k+1} = (y, b)$ with $y = -\frac{1}{2}(a+b)$. In the quadratic coordinates $(A, B, G) = (a^2, b^2, ab)$ each child's form is exactly linear in the parent's, via

$$M_1 = \begin{pmatrix} 1 & 0 & 0 \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \\ -\frac{1}{2} & 0 & -\frac{1}{2} \end{pmatrix} \ [(a,y)], \qquad M_2 = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \\ 0 & 1 & 0 \\ 0 & -\frac{1}{2} & -\frac{1}{2} \end{pmatrix} \ [(y,b)],$$

and $M_1 + M_2$ has left eigenvector $(\alpha, \alpha, 1)$ (the form ϕ) with eigenvalue $\lambda = \frac{1+\sqrt{17}}{4}$, recovering the block identity $\Phi_n = \lambda^n \phi(1, -\frac{1}{3})$. Fix $m \geq 1$, $d = 3^m$, and for each residue $r \in \mathbb{Z}/d$ track the restricted quadratic masses $A_r = \sum_{k \equiv r\,(d)} \rho(k)^2$, $B_r = \sum_{k\equiv r} \rho(k+1)^2$, $G_r = \sum_{k\equiv r} \rho(k)\rho(k+1)$; here A_r is exactly the index-r sublattice ℓ^2-mass. Since 2 is a unit modulo 3^m, $k \equiv r \Rightarrow 2k \equiv 2r$, $2k+1 \equiv 2r+1$, so one substitution step is the exact linear map $v_{n+1} = T_m v_n$ on $\mathbb{C}^3 \otimes \mathbb{C}^{\mathbb{Z}/d}$,

$$T_m = P_1 \otimes M_1 + P_2 \otimes M_2, \qquad (P_1 f)_{r'} = f_{2^{-1}r'},\ (P_2 f)_{r'} = f_{2^{-1}(r'-1)}.$$

Summed over all residues, T_m collapses to $M_1 + M_2$, recovering Theorem 7.4.

Lemma 8.1 (Residue-uniform eigenvector)**.** *Let $t_* = (1, 1, \frac{\sqrt{17}-5}{2})$ be the λ-eigenvector of $M_1 + M_2$. The residue-uniform vector $u_* = (t_*, \ldots, t_*)$ satisfies $T_m u_* = \lambda u_*$ for every $m \geq 1$.*

Proof. If every residue carries the same triple t, then each child residue r' receives $M_1 t$ from its unique M_1-parent $2^{-1}r'$ and $M_2 t$ from its unique M_2-parent $2^{-1}(r'-1)$, i.e. $(M_1+M_2)t$, again residue-uniform. Thus the uniform 3-plane is T_m-invariant and T_m acts on it as $M_1 + M_2$; its top eigenpair (λ, t_*) lifts. □

Because u_* carries the *identical* A-mass 1 in every residue, the top eigenvector is perfectly balanced: provided λ is dominant and simple, $v_n/\lambda^n \to c\, u_*$ and $A_r / \sum_{r'} A_{r'} \to 1/d = 3^{-m}$, i.e. the sublattice mass equidistributes with no anomalous concentration.

Theorem 8.2 (Level-one non-concentration, unconditional)**.** *As $N \to \infty$,*

$$\frac{L(1,N)}{M(N)} = \frac{1}{3} + O\big(N^{-\gamma}\big), \qquad \gamma = \log_2(\sqrt{17} - 3) = 0.1674947\ldots,$$

and the deviation is eventually one-signed (from above). In particular Problem 7.3 holds unconditionally at level one with $\delta = 1$.

Proof. The 9×9 matrix T_1 has the exact spectrum of Theorem 8.3, in which λ is simple and strictly dominant, with residue-balanced eigenvector (Lemma 8.1). Apply the imbalance functional $f(v) = A_0 - \frac{1}{3}(A_0 + A_1 + A_2)$, which annihilates u_*. On the exact initial state,

$$f(T_1^n v_0) = c_2\, \mu_2^n + c_1 \cdot 1^n + c_0(\mu_2')^n + O(1), \qquad \mu_2 = \tfrac{5+\sqrt{17}}{8},\ \mu_2' = \tfrac{5-\sqrt{17}}{8},$$

where μ_2 is the dominant imbalance eigenvalue and its negative partner $-\mu_2$ carries coefficient 0 (hence no oscillation, one sign). Since $\sum_r A_r(n) \sim c\lambda^n$ $(c > 0)$, $A_0/\sum_r A_r - \frac{1}{3} = \Theta((\mu_2/\lambda)^n) = \Theta(N^{-\gamma})$ at $N = 2^n$, with $\lambda/\mu_2 = \sqrt{17} - 3$; monotone interpolation extends the rate to all N, and passing from the masses to L, M shifts by $O(N^{-\beta})$, $\beta > \gamma$. □

Theorem 8.3 (Exact spectrum of the mod-3 refinement)**.** $\operatorname{spec}(T_1) = \big\{\frac{1+\sqrt{17}}{4},\ \pm\frac{5+\sqrt{17}}{8},\ 1\,(\times 2),\ -1,\ \frac{1-\sqrt{17}}{4},\ \pm 5$ *with $\lambda = \frac{1+\sqrt{17}}{4}$ simple and top. For general m, diagonalising the residue variable by the additive characters $\omega = \mathrm{e}(1/3^m)$ gives the exact renewal recursion $V_\xi(n{+}1) = (M_1{+}\omega^\xi M_2)\, V_{2\xi}(n)$, so ξ couples only to 2ξ. The uniform block $\xi = 0$ carries λ; the length-2 orbit $\xi \in \{3^{m-1}, 2 \cdot 3^{m-1}\}$ (cube-root phases $w = \mathrm{e}(\frac{1}{3})$) carries the imbalance radius $\mu_2 = (5 + \sqrt{17})/8$, since $(M_1 + wM_2)(M_1 + w^2 M_2)$ has characteristic polynomial $(x-1)(64x^2 - 84x + 1)/64$ with top root μ_2^2. Hence μ_2 is an imbalance eigenvalue-magnitude of T_m at* every *level.*

All identities are exact in $\mathbb{Q}(\sqrt{17})$; the imbalance radius equals μ_2 (strictly below λ, as $\lambda/\mu_2 = \sqrt{17} - 3 > 1$) for all $m \le 6$, extended to $m \le 10$ in Section 9. The rate $\gamma = \log_2(\sqrt{17} - 3)$ is thus level-independent.

Theorem 8.4 (Per-level non-concentration, $\delta = 1$)**.** *Fix m with $\mathrm{radius}_{\mathrm{imb}}(T_m) < \lambda$ (proved for all $m \le 10$, and on the dominant mode of every m). Then*

$$\frac{L(m,N)}{M(N)} = 3^{-m}\big(1 + O(N^{-\gamma})\big), \qquad \gamma = \log_2(\sqrt{17} - 3),$$

so the limiting per-level ratio is exactly $1/3$: the concentration exponent is $\delta = 1$, with the 3^{-m} exact and carrying no power correction.

Proof. By Lemma 8.1, λ is a simple dominant eigenvalue of T_m with residue-balanced eigenvector, so $v_n = c\lambda^n u_* + O(\mathrm{radius}_{\mathrm{imb}}^n)$ and $A_r/\sum_{r'} A_{r'} = 3^{-m} + O((\mu_2/\lambda)^n) = 3^{-m} + O(N^{-\gamma})$ for each residue; specialise to $r = 0$ and pass from masses to L, M. □

Theorem 8.4 upgrades the measured $\delta \approx 1$ to a theorem for every fixed m, sharpening Theorem 7.5 ($\delta' \approx 0.33$) to the optimal $\delta = 1$ and reinterpreting the earlier "$\approx 0.30 < \frac{1}{3}$": that value is the pre-asymptotic $L(m,N)/L(m-1,N)$, which rises to its true limit $\frac{1}{3}$ from below while the absolute fraction $L(1,N)/M(N)$ falls to $\frac{1}{3}$ from above. The single remaining hypothesis — that no finer character orbit overtakes the cube-root orbit, uniformly in m — is the 2×2 cocycle bound of Section 9.

8.2. **The ergodic-optimization lemma.** The dominance of the cube-root orbit has a clean variational cause, which also pins the resonance constant of Remark 6.6 and supplies the determinant half of the cocycle analysis below.

Lemma 8.5 (Doubling-orbit optimization)**.** *Let* $T : x \mapsto 2x \bmod 1$ *and* $g(x) = 2\sin^2(\pi x)$. *For every* T*-invariant Borel probability measure* μ*,*

$$\int \log\bigl(2\sin^2(\pi x)\bigr)\, d\mu \;\le\; \log\tfrac{3}{2},$$

with equality iff $\mu = \frac{1}{2}\delta_{1/3} + \frac{1}{2}\delta_{2/3}$*. Equivalently, for every period-*P *doubling orbit* O*,* $\prod_{x\in O} 2\sin^2(\pi x) \le (\frac{3}{2})^P$*, with equality iff* $O = \{\frac{1}{3}, \frac{2}{3}\}$*.*

Proof. Put $W(x) = 1 + \frac{1}{5}\cos 2\pi x \in [\frac{4}{5}, \frac{6}{5}]$ and $V = \log W$ (bounded, continuous). With $c = \cos 2\pi x$ one has $2\sin^2 \pi x = 1 - c$ and $\cos 4\pi x = 2c^2 - 1$, so the *exact* polynomial identity

$$\tfrac{3}{2}\, W(2x) - g(x)\, W(x) = \tfrac{3}{2}\bigl(1 + \tfrac{1}{5}(2c^2-1)\bigr) - (1-c)\bigl(1+\tfrac{c}{5}\bigr) = \frac{(2c+1)^2}{5} \;\ge\; 0$$

holds, with equality iff $c = -\frac{1}{2}$, i.e. $x \in \{\frac{1}{3}, \frac{2}{3}\}$. Both sides positive; taking logs gives the sub-action inequality $\phi + V - V\circ T \le \log\frac{3}{2}$ ($\phi = \log g$), equality only on $\{\frac{1}{3}, \frac{2}{3}\}$. Integrating against T-invariant μ (the coboundary $V - V\circ T$ has integral 0) gives the measure form; equality forces support in $\{\frac{1}{3}, \frac{2}{3}\}$. Summing over a periodic orbit (the V-terms telescope) gives the product form. □

The certificate is finite and closed-form: the ansatz $W = 1 + a\cos 2\pi x$ auto-calibrates the 2-cycle for every a, and demanding $c = -\frac{1}{2}$ be the *minimum* of the residual forces $a = \frac{1}{5}$, making it a perfect square. Hence no interval arithmetic is load-bearing. The potential $\phi = \log(1 - \cos 2\pi x)$ lies just outside Bousch's cosine theorem [4] (it has a logarithmic singularity at $x = 0$), yet the maximizer is the periodic (rotation-number $\frac{1}{2}$, Sturmian) orbit, as that theory predicts.

Corollary 8.6 (Cyclotomic form)**.** *For every* $\langle 2\rangle$*-orbit* O *on* $(\mathbb{Z}/q)^\times$*,* q *odd,* $\prod_{u\in O} |1 - \zeta_q^u| \le 3^{|O|/2}$*, with equality iff* $q = 3$*.*

This is the variational form of the level-one resonance $Q_1(\frac{1}{2}) = 9$ of Theorem 6.1: the per-period geometric mean $\frac{3}{2}$ is attained *uniquely* at $\{\frac{1}{3}, \frac{2}{3}\}$ (Remark 6.6, runner-up ≈ 1.422), so no single resonance breaks the 3^t budget. The determinant factor $|1+\omega|^2 = 4\cos^2(\pi x)$ of the cocycle below is a shifted copy of the same optimization, so its "det" half is already solved here.

9. The uniform-in-m cocycle and the tail certificate

The lone hypothesis of Theorem 8.4 is that the imbalance radius of T_m equals μ_2 for *every* m. We reduce it, by three exact steps, to a single rational number per level, and then to one integral inequality — verified to double precision with interval certification pending.

9.1. **Exact $3\times 3 \to 2\times 2$ reduction.** Write $M(\omega) = M_1 + \omega M_2$; the renewal recursion $V_\xi(n+1) = M(\omega^\xi) V_{2\xi}(n)$ makes the growth rate around a $\times 2$-orbit O of length L equal to $\rho(Q_O)^{1/L}$, where $Q_O = \prod_{j<L} M(\omega^{2^j\xi})$ is the ordered product. Elementary and exact: $\operatorname{tr} M(\omega) = \frac{3}{4}(1+\omega)$, $\det M(\omega) = -(1+\omega)^3/8$.

Lemma 9.1 (Twisted-covector splitting)**.** *The row field $u(\omega) = (-1, \omega, 0)$ satisfies the exact functional equation $u(\omega)M(\omega) = u(\omega^2)$. Consequently every orbit product Q_O carries the eigenvalue 1 (left eigenvector $u(\omega_0)$, by telescoping), and on the invariant 2-plane* $\ker u(\omega) = \operatorname{span}\{(\omega, 1, 0), (0, 0, 1)\}$ *the step acts as the explicit cocycle*

$$\tilde{N}(\omega) = \begin{pmatrix} \dfrac{\omega^2 + 4\omega + 1}{4} & \dfrac{1}{2} \\ -\dfrac{\omega(\omega+1)}{2} & -\dfrac{\omega+1}{2} \end{pmatrix}, \qquad \det \tilde{N} = -\tfrac{(1+\omega)^3}{8} = \det M, \quad \operatorname{tr} \tilde{N} = \tfrac{\omega^2+2\omega-1}{4}.$$

Hence $\rho(Q_O) = \max\big(1,\ \rho(\prod_j \tilde{N}(\omega_j))\big)$, and the uniform bound is equivalent *to $\rho(\prod_j \tilde{N}(\omega_j)) \leq \mu_2^L$ for every 3-power orbit. The cube-root 2-cycle attains equality:* $\operatorname{charpoly}(\tilde{N}(w)\tilde{N}(w^2)) = (64x^2 - 84x + 1)/64$*, top root μ_2^2.*

All identities are exact in $\mathbb{Q}(\sqrt{17})$; the $(x-1)$ factor of the TXPOF cube-root polynomial is exactly the split-off eigenvalue 1. The reduction provably cannot be finished by any fixed certificate.

Theorem 9.2 (Fixed-point obstruction)**.** *There is no twisted-Lyapunov / joint-spectral-radius certificate with constant μ_2 over the closed circle, and no finite-degree meromorphic certificate poled at $\omega = 1$ rescues it. The reason is that $\omega = 1$ is a genuine fixed point of $\omega \mapsto \omega^2$ at which $\rho(\tilde{N}(1)) = \lambda > \mu_2$; a submultiplicative certificate would force $\lambda \leq \mu_2$. The bound for $\xi \neq 0$ orbits is therefore a genuine "maximizing orbit that avoids a fixed point" statement: the $O(m)$-step expanding excursion a long orbit makes near $\omega = 1$ is repaid by contraction elsewhere on the same orbit, so only the periodic (spectral) rate — not any finite window — is $\leq \mu_2$.*

(Contrast Lemma 8.5, whose fixed point $x = 0$ has $g(0) = 0$, a local *minimum*; there a bounded sub-action closes the problem. Swapping minimum for maximum at the fixed point is the whole difference, and is why finite windows of g may exceed $\frac{3}{2}$ while invariant averages cannot.)

9.2. **The primitive-orbit trace.** For a primitive 3^m-th root ω_0 the $\times 2$-orbit $\omega_j = \omega_0^{2^j}$ visits $e(2\pi u/3^m)$ over *all* units u modulo 3^m exactly once (as 2 is a primitive root), period $P = 2 \cdot 3^{m-1} = \varphi(3^m)$; there is exactly one such orbit. Put

$$Q_m = \prod_{j<P} \tilde{N}(\omega_j), \qquad T_m = \operatorname{tr} Q_m, \qquad D_m = \det Q_m.$$

For $m \geq 2$ these primitive orbits are exactly the non-cube-root 3-power orbits (the $m = 1$ orbit is the excluded cube-root 2-cycle), so the goal is $\rho(Q_m) < 1$ for all $m \geq 2$.

Theorem 9.3 (Collapse to a scalar)**.** (i) Determinant. *$D_m = 8^{-P}$ exactly. Indeed $D_m = (-\frac{1}{8})^P \big(\prod_u (1+\zeta^u)\big)^3$ and $\prod_u (1+\zeta^u) = (-1)^P \Phi_{3^m}(-1) = 1$, since $\Phi_{3^m}(x) = x^{2\cdot 3^{m-1}} + x^{3^{m-1}} + 1$ gives $\Phi_{3^m}(-1) = 1 - 1 + 1 = 1$ and P is even.*

(ii) Rationality. *$T_m \in \mathbb{Q}$. As $\operatorname{Gal}(\mathbb{Q}(\zeta_{3^m})/\mathbb{Q}) = (\mathbb{Z}/3^m)^\times = \langle 2 \rangle$, an automorphism $\zeta \mapsto \zeta^{2^k}$ cyclically rotates the ordered product $\prod_j \tilde{N}(\omega_j)$, and the trace is rotation-invariant. (Empirically $T_m = N_m/4^P$, N_m odd.)*

(iii) Schur–Cohn. *Q_m has characteristic polynomial $x^2 - T_m x + 8^{-P}$, so $\rho(Q_m) < 1 \iff |T_m| < 1 + 8^{-P}$; the discriminant $T_m^2 - 4 \cdot 8^{-P}$ is positive for every m, so the spectrum* *is real positive and* $\rho(Q_m) < T_m$. *Therefore the uniform* $\delta = 1$ *step is equivalent to the scalar statement*

$$T_m < 1 \quad \textit{for all } m.$$

The values (exact for $m \le 5$, high precision through $m \le 10$):

m	P	T_m	rate $\rho(Q_m)^{1/P}$
2	6	$2193/4^6 = 0.5354004$	0.9011121
3	18	$37057403985/4^{18} = 0.5392562$	0.9662727
4	54	0.3802907	0.9822553
5	162	0.04344032	0.9808259
6	486	5.87575×10^{-5}	0.9801541
7–10	≥ 1458	$\le 2.03 \times 10^{-13}$	$\to 0.9801700$

Thus $T_m \downarrow 0$ super-exponentially with $\max_m T_m = T_3 \approx 0.5393$ and margin $1 - T_m \ge 0.46$ throughout; the top Lyapunov exponent is $\approx -0.0200 < 0$ (the two exponents sum to $\int \log|\det \tilde{N}| = -\log 8$, since $\int \log|1+\omega| = 0$), so the equidistributed limit contracts. No fixed submultiplicative norm can supply $T_m < 1$ — the same $\omega = 1$ obstruction (avg $\log \|\tilde{N}\|_F = +0.314 > 0$) — so $|T_m| < 1$ is a genuine renewal/excursion statement, invisible to any per-step norm.

9.3. **The tail closure.** The final step evades the obstruction *arithmetically.* Set $\varphi_b(\theta) = \frac{1}{b}\log\|\mathrm{block}_b(\theta)\|_F$ with $\mathrm{block}_b(\theta) = \tilde{N}(\mathrm{e}(\theta))\tilde{N}(\mathrm{e}(2\theta))\cdots\tilde{N}(\mathrm{e}(2^{b-1}\theta))$.

Lemma 9.4 (Block-tiling average, Frobenius)**.** *For every* b *and* m*,*

$$\log\sigma_1(Q_m) \;\le\; \log\|Q_m\|_F \;\le\; \frac{1}{b}\sum_{j<P}\log\|\mathrm{block}_b(\theta_j)\|_F \;=\; \sum_{u \ \mathrm{unit}\,(3^m)} \varphi_b(u/3^m).$$

Proof. Frobenius is submultiplicative, so for each of the b tilings of the cyclic orbit into consecutive b-blocks, $\log\|Q_m\|_F \le \sum_{\text{block starts}} \log\|\mathrm{block}_b\|_F$. Averaging the b tilings, each orbit index is a block-start exactly once. ☐

Lemma 9.5 (Exact Ramanujan equidistribution)**.** *For a trigonometric polynomial* h *of degree* D *and every* m *with* $3^{m-1} > D$*,*

$$\frac{1}{\varphi(3^m)}\sum_{u \ \mathrm{unit}\,(3^m)} h(u/3^m) = \int_0^1 h(\theta)\,d\theta \qquad \textit{exactly.}$$

Proof. $\sum_{u \ \mathrm{unit}\,(3^m)} \mathrm{e}(ku/3^m) = c_{3^m}(k)$, a Ramanujan sum, which vanishes unless $3^{m-1} \mid k$; for $|k| \le D < 3^{m-1}$ only $k = 0$ survives, giving the mean $\hat{h}(0) = \int h$. ☐

The units $u/3^m$ *avoid* $\{0, \frac{1}{3}, \frac{2}{3}\}$ (nearest approach 3^{-m}), so Lemma 9.5 never samples the positive spikes of φ_b at the fixed point ($\omega = 1$, rate $\log\lambda$) or the cube-root cycle ($\log\mu_2$): it requires φ_b only to have a negative *average*, not to be pointwise ≤ 0. This is precisely how the $\omega = 1$ obstruction of Theorem 9.2 — which forces any pointwise sub-action to be $\ge \log\lambda > 0$ at $\omega = 1$ — is bypassed.

Proposition 9.6 (Certificate chain)**.** *If* $B \ge \varphi_b$ *is a trig-polynomial majorant with* $\int_0^1 B < 0$ *and* $\deg B < 3^{m-1}$*, then*

$$\log\sigma_1(Q_m) \;\le\; \sum_u \varphi_b(u/3^m) \;\le\; \sum_u B(u/3^m) \;=\; P\int_0^1 B \;<\; 0,$$

so $\sigma_1(Q_m) < 1$ and $T_m < 1$. Such a B exists because $\int_0^1 \varphi_b < 0$ for $b \geq 14$. Concretely,

$$T_m < 1 \impliedby \sum_{u \text{ unit}(3^m)} \varphi_{16}(u/3^m) < 0 \impliedby \int_0^1 \varphi_{16} < 0.$$

Numerically $\int_0^1 \varphi_{16} = -0.00465 \pm 0.00022$ ($> 20\sigma$), and the exact per-m units-average $\frac{1}{P}\sum_u \varphi_{16}(u/3^m)$ equals -0.004743 to within 2×10^{-5} for every $m \geq 6$; the resulting rigorous bounds on $\log \sigma_1(Q_m)$ run from -2.3 ($m = 6$) to -5042 ($m = 13$). With $b = 16$ the frequency content of φ_b reaches $\sim 2^{18}$, giving worst-case threshold $m_0 = 13$; the seam closes with no gap.

Remark 9.7 (Honest status)**.** Lemmas 9.4–9.5 and Proposition 9.6 are exact, and the reduction of the uniform-$\delta = 1$ step to $\int_0^1 \varphi_{16} < 0$ is rigorous and structurally bypasses the $\omega = 1$ obstruction. The residual is a *finite deterministic computation*, not a conceptual gap: (a) interval-arithmetic certification of the finite sums $\sum_u \varphi_{16}(u/3^m)$ (and of $\int \varphi_{16}$) — routine `mpmath` work, margins from 10^0 to 10^{-2187}, not yet executed as intervals; and (b) construction of one explicit trig-polynomial majorant B. Coverage is complete: $m \leq 10$ (exact/high-precision T_m, Theorem 9.3), $m = 11, 12, 13$ (per-m certificate), $m \geq m_0 = 13$ (uniform majorant). We cite the interval certification as the final computational step; it is *not* claimed as executed here.

9.4. **Status of the dependency chain.** Assembling Sections 8–9, the program's chain now reads

$$[\,T_m < 1\ \forall m\,] \implies (\text{TM-}\ell^2,\ \delta = 1 \text{ uniform}) \implies \text{K2} \implies (\text{MS},\ \tfrac{2}{3}\text{ exact}) \overset{+(\text{MULT-mass})}{\implies} c_0 = \tfrac{1}{2} - \varepsilon.$$

Every arrow is now established except the two named inputs. **As of v6 the first is discharged:** the certification run of §13 closes $T_m < 1$ for all $m \geq 2$, so uniform $\delta = 1$ is unconditional and (MS) is a theorem. The three analytic lemmas beneath it are moreover machine-verified (see §11). The second, (MULT-mass) (Conjecture 6.9), is measured favorable but not proved, and is now the *only* remaining input to the sharp constant. Accordingly:

- The *unconditional* depth constant is unchanged at $c_0 \approx 0.124$ (Theorem 1.2).
- The sharp $c_0 = \frac{1}{2} - \varepsilon$ (Theorem 6.10) now rests on exactly *one* input, (MULT-mass), with everything else on the $\delta = 1$ route proved. We do *not* claim $c_0 = \frac{1}{2}$ unconditionally: the K2 assembly details and (MULT-mass) remain to be written out as theorems before the sharp constant can be asserted rigorously.

The four previously independent attacks on the residual non-concentration (signed NAF/Ramanujan sums, Thue–Morse lattice masses, the finite transfer spectra of Theorem 6.5, and the hot-window/cocycle route) have now provably converged onto this one certifiable scalar. T2 (the quenched Bernstein–Lagarias wall) is untouched and remains honestly out of scope.

10. An exact constant for the Thue–Morse carry operator

The triadic eigen-form of §8 was introduced to obtain sharp per-level non-concentration. It has an unrelated consequence worth recording, because it replaces a quantitative estimate in the literature by an exact algebraic number.

Recent work on Gowers-norm bounds for digital sequences proceeds through a carry-state transfer matrix: states are carry vectors, transitions are induced by the base-b addition $\delta(r; e)_w = \lfloor (r_w + \langle 1, w \rangle \cdot e)/q \rfloor$, and the argument extracts a contraction from an ℓ^∞ gap for $|M|$, established by path coupling rather than spectrally (the transition weights are complex for general (q, b)). For $q = b = 2$ — the Thue–Morse case — those weights are real, and the matrix may be diagonalised.

Proposition 10.1. *For $q = b = 2$ the carry-state transfer matrix M_k has spectral radius*

$$\rho(M_k) \;=\; \frac{1+\sqrt{17}}{8},$$

independently of k. Consequently the associated exponent satisfies

$$\eta_0 \;=\; 1 - \log_2\Big(\frac{1+\sqrt{17}}{4}\Big) \;=\; 0.6429813631394237\ldots \;=\; 1-\beta,$$

where $\beta = \log_2((1+\sqrt{17})/4)$ is the energy exponent of Theorem 7.4.

Status of proof. At $k = 2$ the characteristic polynomial factors exactly, the relevant factor being $4\lambda^2 - \lambda - 1$, whose positive root is $(1+\sqrt{17})/8$. The identification has been confirmed numerically to 10^{-16} at $k = 3$ and to 7 decimal places at $k = 4$. A uniform-in-k proof — the natural route being the same substitution identity that yields the eigen-form $\varphi = \alpha(a^2 + b^2) + ab$, $\alpha = (5 + \sqrt{17})/4$, of §8 — is not written out here, and the statement should be read as a verified computation rather than as a theorem of this paper. □

Two things make this worth stating even in that provisional form. First, the constant is *k-independent*, whereas the published explicit bounds decay doubly exponentially in k (at $k = 3$ the published value is $\approx 1.31 \times 10^{-13}$, smaller by a factor of about 5×10^{12}). Second, and more to the point of this paper, η_0 and β are the same algebraic number: the carry operator governing digital Gowers norms and the spectral measure governing Thue–Morse Fourier energy are two faces of one object. That was not visible from either side alone, and it is the reason §8 generalises beyond the use made of it here.

11. Machine verification of the $\delta = 1$ certificate

The lemmas underlying §9 have been formalised in Lean 4 against Mathlib. We record the scope precisely, since the value of such an exercise lies entirely in what it does *not* cover.
What is verified. Three groups of results, 79 declarations in total:

- the Schur–Cohn stability criterion for a real 2×2 companion, in the scalar form used by the cocycle reduction;
- Ramanujan-sum exactness for 3^m: the character sum over $(\mathbb{Z}/3^m)^\times$ vanishes for every nonzero frequency of modulus below 3^{m-1}, which is Lemma B of the tail closure;
- the Jackson-kernel results: the normalisation $\int_0^1 (\sin n\pi t/\sin \pi t)^4 dt = n(2n^2+1)/3$ and the weighted bound $\int_{-1/2}^{1/2} J_n(t)|t|\,dt \le 3/(4n)$;
- a kernel-only certified rational interval arithmetic, with soundness proved for the arithmetic operations and for Horner evaluation, for use on the finite residual.

The trust boundary. Every declaration passes `#print axioms` depending only on `propext`, `Classical.choice` and `Quot.sound`. No declaration depends on `sorryAx`, and the development uses neither `decide` nor `native_decide`, so the trusted base is not enlarged beyond Lean's kernel and Mathlib. This check is run on every build rather than asserted: searching the sources for admitted statements would not suffice, since a declaration can inherit one through an import without the token appearing in its own file.

What is not verified. The reduction chain of §9.4 is *not* formalised — only the analytic lemmas at its base. Neither is (MULT-mass), nor the certification run of §13, whose remaining obstacle is certified interval arithmetic for transcendental functions at a scale the kernel cannot replay. Nothing in this section bears on the Collatz conjecture.

Remarks on the exercise. Mathlib contains no Fejér or Jackson kernel and no Ramanujan sums, so both developments were built from primitives. The route for the normalisation is worth recording as it is shorter than the classical one: expanding $\|\sum_{j<n} e(jt)\|^4$ over quadruples and integrating termwise by character orthogonality converts the integral into the cardinality $\#\{(i,j,k,l) \in [0,n)^4 : i+j = k+l\}$, which fibres over $d = i-k$ into $\sum_{|d|<n}(n-|d|)^2 = n(2n^2+1)/3$. The weighted bound then follows from two pointwise estimates split at $|t| = 1/(2n)$; that constant is forced, being the minimiser of $c^2/2+1/(32c^2)$, and the natural choice $1/n$ does not suffice.

12. Open problems and the quenched wall

Open Problem 12.1 (The kernel inequality)**.** Prove Conjecture 6.4: $B_t(r) \le C\,3^t$ in the mid-window. By Theorem 6.3 this is one-sided and per-level; by the Thue–Morse identification the infinite-window version concerns Fourier coefficients of a classical singular measure [14], and the residual arithmetic input is the signed-weight sum bound for multiples of powers of 3 stated above. Caution: the unsigned weighted representation count is $\Theta(1)$ per integer, so minimal-weight bounds alone do not suffice. The sharp route of Sections 8–9 supersedes this problem: the triadic eigen-form proves the governing non-concentration (TM-ℓ^2) at $\delta = 1$ per level, from which K2 follows, modulo only the uniform-in-m scalar $T_m < 1$ (Proposition 9.6, interval certification pending).

Open Problem 12.2 (T2: the annealed–quenched jump)**.** The quenched statement (a single orbit, k up to the orbit length $\approx 2.41\log_2 n_0$) requires a k-uniform seed×depth type-II bilinear bound $\sum_{n\sim N}\sum_{k\sim K} a_n b_k\,\mathrm{e}(\alpha\,\mathrm{s}_2(m_k(n))) = o(NK)$, which fails at cylinder aggregation because $s \mapsto c(s)$ is non-arithmetic across words: this is the Bernstein–Lagarias obstruction [3] in exponential-sum form. We regard T2 as the genuinely Collatz-hard boundary; nothing in this paper crosses it.

13. Numerical results

All computations are exact (big-integer orbits; machine-precision products with error tracking). Highlights, with data files listed in the reproducibility note:

Quantity	Measured	Interpretation
$\lvert S_k(\alpha; N)\rvert/N$, spikes	only at $\alpha \in \mathbb{Z}$	direct evidence for Theorem 1.1
breakdown depth $k^*(N)$	$2.567 \log_2 N$, char.-independent	orbit absorption $1/\log_2 \frac{4}{3} = 2.409$, not c_0
onset length $L^*(d)$, $d = 3 \dots 3^{20}$	≈ 200–400, flat	empirical $\theta = 0$, $c_0^{\mathrm{emp}} = \frac{1}{2}$
$\sum_a \lvert S_a\rvert^2/N$ at $\alpha = \frac{1}{3}$	$O(1)$	square-root cancellation in mean square
$\sum_a \lvert S_a\rvert^2/N$ at $\alpha = \frac{1}{2}$	$\sim N^{2\gamma-1}/d$	mass on the two Newman–Coquet freque
$Q_m(\frac{1}{2})$, $m \le 7$	9 (to 2×10^{-12})	Theorem 6.1
deep-window sums	$= 3^m 2^r$ (to 8×10^{-16})	Theorem 6.2
$B_1(r)$, $r \le 40$	$= 2(\frac{3}{2})^r$ (to 7×10^{-15})	Theorem 6.3(*i*)
$B_t/(2 \cdot 3^{t-1})$, $t \le 14$	$\le 1.04 \cdot 3^t/P_t$, no trend	Conjecture 6.4, Thm 6.5
$r_t = 9^{1/P_t}/2$, $t \le 6$	≤ 0.7211, $\downarrow \frac{1}{2}$ (to 10^{-6})	Theorem 6.5
$B_t = 2^{-r}[3^t S_t - 3^{t-1} S_{t-1}]$	rel. err $\le 10^{-14}$	Theorem 7.1
$M(N)/N^\beta$, $N \le 10^7$	$\to$ const, $\ge \frac{2}{9}$	Theorem 7.4 ($\beta = 0.35702$)
cycle census, period ≤ 20	$\max \mathrm{gm} = \frac{3}{2}$ unique	Remark 6.6
κ_j (mass collision), $k = 8, 10, 12$	1.4–4.8 (mode class)	Conjecture 6.9 (MULT-mass)
unsigned rep.-count, odd W	$\to 1$ (0.97–1.00)	minimal-weight route excluded
kernel $X_m(r)$, $m = 5, 6, 7$	$\frac{2}{3} \cdot 2^{(2\gamma-1)r}/3^{m-1}$	constant now exact (Thm 6.3)
offset support	exactly $2d/3$	$3 \nmid a(s)$, proved
$L(m, N)3^m/M(N)$, $m \le 4$	$\to 1$ (from above/below)	Theorems 8.2, 8.4 ($\delta = 1$)
imbalance radius of T_m, $m \le 10$	$= (5 + \sqrt{17})/8$ exactly	Theorem 8.3; level-independent
$\max_\mu \int \log(2 \sin^2 \pi x)$	$= \log \frac{3}{2}$, unique $\{\frac{1}{3}, \frac{2}{3}\}$	Lemma 8.5
$D_m = \det Q_m$, $m \le 7$	$= 8^{-P}$ (to 10^{-13})	Theorem 9.3(*i*), $\Phi_{3^m}(-1) = 1$
T_m, $m = 2 \dots 10$	≤ 0.5393, $\downarrow 0$, margin ≥ 0.46	Theorem 9.3(*iii*)
Lyapunov exp. of $\tilde{N}$-cocycle	$\approx -0.0200 < 0$	sum of exps $= -\log 8$
$\int_0^1 \varphi_{16}$ / per-m $\sum_u \varphi_{16}$	-0.0047 ($> 20\sigma$) / < 0, $m \le 13$	Prop. 9.6 (interval cert. pending)

Reproducibility. Verification scripts and data: `txpof18_expsum_probe.py`, `txpof20_cancellation_surface` `txpof21_theta_probe.py`, `txpof22_T1_verify.py`, `txpof23_ms_probe.py`, `txpof24_transfer_sum.py`, `txpof35_checks.py`, `txpof36_spectra.py`, `txpof37_mult_perclass.py`, `txpof38_hotwindow.py`, `txpof39_energy_lb.py`, `txpof40_multmass.py`, `txpof42_triadic_eigenform.py`, `txpof43_cyclotomic_opt.p` `txpof44_cocycle.py`, `txpof45_trace.py`, `txpof46_tail.py` and associated CSV tables (`txpof42_spectrum.cs` `txpof44_certificate.csv`, `txpof45_traces.csv`, `txpof46_certificate.csv`). The energy bound (Theorem 7.4), the transfer spectrum (Theorems 8.3, 9.3), and the cocycle reduction (Lemma 9.1) are verified exactly in $\mathbb{Q}(\sqrt{17})$ and $\mathbb{Q}(\zeta_{3^m})$; the tail sums of Proposition 9.6 are double-precision with interval certification pending (Remark 9.7).

Acknowledgements. This paper consolidates a ticketed research program (TXPOF-18, -20, -21, -22, -23, -24, -35, -36, -37, -38, -39, -40, -42, -43, -44, -45, -46); intermediate corrections are recorded in Remarks 4.1, 6.8 and 6.6.

References

[1] M. Baake and U. Grimm, *The singular continuous diffraction measure of the Thue–Morse chain*, J. Phys. A: Math. Theor. **41** (2008), 422001. arXiv:0809.0580.
[2] M. Baake, F. Gähler and U. Grimm, *Scaling of the Thue–Morse diffraction measure*, arXiv:1311.4371 (2013).
[3] D. Bernstein and J. C. Lagarias, *The 3x+1 conjugacy map*, Canad. J. Math. **48** (1996), 1154–1169.
[4] T. Bousch, *Le poisson n'a pas d'arêtes*, Ann. Inst. H. Poincaré Probab. Statist. **36** (2000), 489–508.

[5] J. Coquet, *A summation formula related to the binary digits*, Invent. Math. **73** (1983), 107–115.
[6] J.-M. Deshouillers, L. Habsieger, S. Laishram and B. Landreau, *Sums of the digits in bases 2 and 3*, arXiv:1611.08180 (2016).
[7] É. Fouvry and C. Mauduit, *Sommes des chiffres et nombres presque premiers*, Math. Ann. **305** (1996), 571–599.
[8] A. O. Gelfond, *Sur les nombres qui ont des propriétés additives et multiplicatives données*, Acta Arith. **13** (1967/68), 259–265.
[9] C. Hercher, *There are no Collatz m-cycles with* $m \leq 91$, J. Integer Seq. **26** (2023), Article 23.3.5.
[10] D.-H. Kim, *On the joint distribution of q-additive functions in residue classes*, J. Number Theory **74** (1999), 307–336.
[11] C. Mauduit and J. Rivat, *Sur un problème de Gelfond: la somme des chiffres des nombres premiers*, Ann. of Math. **171** (2010), 1591–1646.
[12] C. Müllner and L. Spiegelhofer, *Normality of the Thue–Morse sequence along Piatetski-Shapiro sequences, II*, Israel J. Math. **220** (2017), 691–738.
[13] D. J. Newman, *On the number of binary digits in a multiple of three*, Proc. Amer. Math. Soc. **21** (1969), 719–721.
[14] M. Queffélec, *Substitution Dynamical Systems — Spectral Analysis*, Lecture Notes in Mathematics **1294**, Springer, 2nd ed., 2010.
[15] H. G. Senge and E. G. Straus, *PV-numbers and sets of multiplicity*, Period. Math. Hungar. **3** (1973), 93–100.
[16] L. Spiegelhofer, *The level of distribution of the Thue–Morse sequence*, Compos. Math. **156** (2020), 2560–2587.
[17] L. Spiegelhofer, *Collisions of digit sums in bases 2 and 3*, Israel J. Math. (2023). arXiv:2105.11173.
[18] C. L. Stewart, *On the representation of an integer in two different bases*, J. Reine Angew. Math. **319** (1980), 63–72.
[19] T. Stérin and D. Woods, *The Collatz process embeds a base conversion algorithm*, arXiv:2007.06979 (2020).
[20] T. Tao, *Almost all orbits of the Collatz map attain almost bounded values*, Forum Math. Pi **10** (2022), e12.
[21] T. Tao, *The Collatz conjecture, Littlewood–Offord theory, and powers of 2 and 3*, blog post, 2020-01-25.