mathdb:2611.00026v1 04 Nov 2026

A convergent cascade series without geometric decay

Research manuscript

29 October 2026
Independent review and proof replay completed; no priority claim

Abstract

We give an explicit binary saturated context tree, with exactly one infinite branch and no infinite shift-stable subtree, equipped with fair probabilities at every finite and infinite context. Its cascade series for the fixed alpha-LIS 1 converges with sum at most 2, but admits no eventual bound Cr^n with $C > 0$ and $0 < r < 1$. The index is the length grading of equation (5) of Cénac, Chauvin, Noûs, Paccaut and Pouyanne [1]. This contradicts the explicit geometric-rate clause of their Conjecture 4.2, under this literal rate interpretation. It does not refute the universal bare-convergence conjecture. Classical prefix-code counting and Kraft's inequality receive no novelty credit. A closed Lean certificate accompanies the construction; formal checking supplies no literature-priority conclusion.

1 Source conventions and exact claim

A word is a finite string over $\{0, 1\}$; ϵ is empty. Prefixes and suffixes retain the written left-to-right order. A rooted context tree T is prefix closed and saturated: every internal node has both children. Finite contexts are leaves; infinite contexts are branches all of whose finite prefixes are nodes. We require at most countably many infinite contexts and a nontrivial tree. A probabilisation assigns to each finite *and infinite* context c a positive law $q_c(0) + q_c(1) = 1$.

For a nonempty noninternal word w, its LIS s is its longest internal *strict* suffix. Write $w = \beta_1 \cdots \beta_p \alpha s$. The alpha-LIS is αs. For any noninternal word v, let $\mathrm{cont}(v)$ be its unique context prefix. With σ deleting the first letter, the source cascade is

$$\mathrm{casc}(w) = \prod_{i=1}^{p} q_{\mathrm{cont}(\sigma^i w)}(\beta_i).$$

The minimality implicit in LIS makes each context in this product well defined. For a context alpha-LIS t, the source series has terms

$$\kappa_t(n) = \sum_{\substack{c \text{ finite context: } \alpha\text{-LIS}(c)=t \\ |c|=|t|+n-1}} \mathrm{casc}(c), \qquad n \geq 1.$$

The certificate also sets the $n = 0$ term, which is zero for $t = 1$ here.

2 Statement and tree construction

A stable tree has $\sigma(T) \subseteq T$. By an infinite shift-stable subtree we mean an infinite prefix-closed node set contained in T and closed under σ; our exclusion below even allows subtrees without saturation. Define geometric decay to mean

$$\exists C > 0\ \exists r \in (0,1)\ \exists n_0\ \forall n \geq n_0, \quad \kappa_t(n) \leq Cr^n.$$

The source's phrase "geometrical rates" is not a separate formal definition. We use this standard length-based meaning, tied to its equation (5). An exponential bound on nonnegative series tails implies this term bound, so that stronger interpretation is excluded too.

Theorem 1. *There is a non-null probabilised binary context tree with countably many contexts and no infinite shift-stable subtree, for which* 1 *is a context alpha-LIS and*

$$\sum_{n\geq 0} \kappa_1(n) \leq 2, \qquad \neg\exists C > 0,\ r \in (0,1),\ n_0 :\ \forall n \geq n_0, \kappa_1(n) \leq Cr^n.$$

Let $b = (01)^\infty$ and let p_m be its prefix of length m. Put

$$N_0 = 0, \quad N_k = k^2(N_{k-1} + k + 10), \qquad a_k = (01)^k 00 \quad (k \geq 1).$$

Set

$$I = \{p_m : m \geq 0\} \cup \bigcup_{k\geq 1} \{a_k v : |a_k v| < N_k\}, \quad T = I \cup \{wa : w \in I,\ a \in \{0,1\}\}.$$

The a_k are prefix incomparable: for $j < k$, their letters at position $2j+2$ (numbering from one) disagree. Each a_k is the off-spine child of p_{2k+1}. The recurrence gives strictly increasing depths and $N_k > 2k+2$; in particular $N_1 = 11$ and $N_2 = 92$. The set I is prefix closed. Completing it by both children gives precisely I as internal set: an internal descendant of a node outside I would contradict prefix closure of I. Thus T is saturated. Each bush has all its bottom leaves at depth N_k.

A branch leaving b either reaches an ordinary leaf immediately or enters one finite bush. Consequently b is the unique infinite context. Finite words are countable, so all contexts are countable. Define $q_c(0) = q_c(1) = 1/2$ on every context, including b. This is normalized and uniformly positive. It represents a fair iid process; no new process-mixing property is claimed.

mathdb:2611.00026v1 04 Nov 2026

3 Excluding stable subtrees and fixing one alpha-LIS

Every word of T of length at least three starts with 01. Indeed spine nodes and all bushes do, and the only shorter off-spine possibilities terminate by depth two. In particular $10, 11 \notin T$. If a prefix-closed, shift-closed $S \subseteq T$ contained a word of length at least three, its three-letter prefix $01a$ would belong to S. Shifting gives $1a \in S \subseteq T$, a contradiction. Such S therefore has only words of length at most two, hence is finite.

Fix $k \geq 3$ and put

$$H = N_{k-1} + 3, \quad R = N_k - (2k+2) - H = kB + r, \quad 0 \leq r < k.$$

Here $R \geq 0$: using $N_k = k^2(N_{k-1} + k + 10)$, subtracting $N_{k-1} + 2k + 5$ leaves a positive integer for $k \geq 3$. For each independent choice of words u_i of length $k-2$, take

$$w = a_k(11u_1)(11u_2)\cdots(11u_B)1^{r+H}.$$

A block $11u_i$ has exactly k letters; there is no trailing extra letter. These words have length N_k and are leaves in bush k. Fixed block boundaries recover every free bit, so their number is exactly $2^{(k-2)B}$.

Lemma 1. *Every nonempty proper suffix of every selected w is noninternal. Consequently its LIS is ϵ and its alpha-LIS is 1.*

Proof. The terminal run has at least $H \geq 3$ ones. Suffixes of length one and two are $1, 11$, neither internal. A longer suffix cannot be a spine prefix, since an alternating spine prefix cannot end in 111.

An internal suffix in bush $j < k$ would have length less than $N_j \leq N_{k-1}$. It would lie wholly in the terminal ones, and could not begin with a_j, which starts with zero.

For $j \geq k$, an anchor occurrence requires an initial alternating segment $(01)^k$ of length $2k$. No such segment starts in the tail after a_k. To see this explicitly, the forced 11 pairs begin k positions apart. A start inside a block before the next block boundary encounters the next complete pair within the next $k+1$ letters. If it is in the last block and no next block exists, it encounters two terminal ones within the next $k+1$ letters. A start at a boundary already encounters 11, and a start in the terminal run does as well. A $2k$ window, whenever it fits, thus contains equal adjacent letters. This also covers $B = 0$, when the whole tail is ones.

mathdb:2611.00026v1 04 Nov 2026

Completion of the suffix proof. A proper start within a_k at a letter 1 cannot start an anchor. A start at even offset $2h$, $1 \le h < k$, leaves only $2(k-h)$ alternating letters before the 00 defect. A start at its penultimate zero meets that defect at once. A start at its last zero meets a tail beginning 11 and alternates for at most two letters. None admits the required $2k$ letters. At offset zero, $j > k$ is blocked by the same 00 defect; $j = k$ gives w itself rather than a proper suffix. These possibilities exhaust both kinds of internal nodes. The empty suffix is internal, so the LIS is empty and the preceding letter is the final 1. □

Also 1 itself is a leaf with alpha-LIS 1, so the chosen index belongs to the source's set of context alpha-LIS.

4 Counting and failure of geometric decay

Fair probabilities give one factor $1/2$ for each letter preceding the alpha-LIS. Thus a selected leaf has cascade $2^{-(N_k-1)}$. Since $|1| = 1$, the equation (5) index is exactly its length, not an arbitrary enumeration of contexts. Define

$$F_k = (k-2)B, \qquad D_k = N_k - 1 - F_k.$$

The selected leaves give

$$\kappa_1(N_k) \ge 2^{F_k-(N_k-1)} = 2^{-D_k}.$$

We have $F_k \le N_k - 1$, so $D_k \ge 0$. Euclidean division gives $B \ge R/k - 1$ and hence

$$\begin{aligned} 0 \le \frac{D_k}{N_k} &\le \frac{2}{k} + \frac{N_{k-1}}{N_k} + \frac{3k+5}{N_k} \\ &\le \frac{2}{k} + \frac{1}{k^2} + \frac{3k+5}{k^2(k+10)} \longrightarrow 0. \end{aligned}$$

For the first inequality use $N_k - 1 - (k-2)B \le 2N_k/k + (2k+2) + H + k$. The recurrence gives the other bounds, as well as $N_k \to \infty$.

Suppose an eventual Cr^n bound existed. The subsequence lower bound would imply

$$-\frac{D_k}{N_k}\log 2 \le \frac{\log C}{N_k} + \log r$$

for all sufficiently large k. The left side tends to zero and the right side to $\log r < 0$, a contradiction. This proves the rate obstruction with all constants and cutoffs quantified. We do not need an optimal rate or an exact root-limit statement.

mathdb:2611.00026v1 04 Nov 2026

mathdb:2611.00026v1 04 Nov 2026

5 Convergence and literal formal scope

Finite contexts form a prefix-free family. For any finite subfamily, the fair binary cylinder probabilities $2^{-|c|}$ sum to at most one, because the cylinders are disjoint. Taking the supremum over finite subfamilies gives the classical countable Kraft bound. In the subfamily with alpha-LIS 1, the cascade is $2^{1-|c|}$. Grouping nonnegative terms by their lengths therefore gives

$$\sum_{n\geq 0} \kappa_1(n) = \sum_{\alpha\text{-LIS}(c)=1} 2^{1-|c|} \leq 2.$$

This completes the theorem. Kraft and its application here are classical helpers, not additional proposed outcomes.

The nine immutable mathematical Lean modules are DepthDefs, SourceDefs, TreeConstruction, DepthArithmetic, DepthAnalysis, SelectedFamily, SelectedCounting, KraftSummability and CascadeTheorem. Booleans encode $0, 1$ as false,true. SourceDefs defines the actual context prefix, LIS, alpha-LIS, probability laws on both context kinds, finite cascade product, and the exact finite length sum. Its default definitions outside valid noninternal inputs are irrelevant: the factor-validity lemmas prove all consumed suffix inputs valid. TreeConstruction supplies the admissible source tree and exclusion of stable subtrees. The final closed endpoints, in namespace `CascadeRate`, are

```
source_geometric_rate_refuted :
  ¬ SourceGeometricRateClaim

convergent_non_geometric :
  Summable (kappa sourceTree (fairAssignment sourceTree) [true])
  ∧ (tsum of that same series ≤ 2)
  ∧ ¬ GeometricDecay (that same series).
```

The accompanying inventory prints actual complete types, including generated, private and instance declarations, and collects transitive axioms. The only allowed axioms are propext, Classical.choice and Quot.sound. Compilation uses Lean 4.34.1 and the pinned eight-package environment. There are no residual source-existence, counting, suffix, convergence or asymptotic hypotheses in these endpoints. The formal universal rate claim quantifies saturated nontrivial source subtrees; the ordinary proof above excludes even the broader node-subtree class. No theorem about all nonfair probabilisations' bare convergence is claimed.

6 Attribution and remaining objections

The target is [1], Definitions 2.5–2.12, equation (5), Proposition 3.1, Definition 3.2 and Conjecture 4.2, printed page 2037. The corresponding preprint is [2], Conjecture 5.2, page 39. The actual 20-page publication-hosted supplement [3] was recovered in the supporting source review through a public archival capture; its live publisher route had returned a challenge. That route was not bypassed or retried for this packet.

The supplement's stable comb examples already allow slow cascades; they do not satisfy absence of an infinite stable subtree. Similarly, classical prefix codes prescribe lengths, but do not by themselves supply both a single fixed alpha-LIS and the exclusion of stable subtrees. The proposed content is this simultaneous realization. Deep finite bushes, Kraft counting, fair iid stationarity and auxiliary arithmetic are explicitly uncredited.

The strongest remaining substantive objection is subsumption: a pre-existing construction, or short explicit implication from one, supplying these same two bridges would defeat discovery credit. The supporting reviews compare classical coding, stable comb and deep-bush examples, but are bounded literature screens rather than global novelty certificates. Their registry observations were refreshed on 29 October 2026 at 08:10 UTC using complete registry trees and relevant primary bodies. No exact MathDB identifier or worldwide priority is asserted. Exact current MathDB solution bodies and an exhaustive specialist citation-database audit remain outside coverage.

The strongest mathematical falsifiers are a selected leaf with a nonempty internal proper suffix, or an infinite shift-stable subtree. The proof above addresses all indices rather than a finite numerical sample. A source clarification assigning a different rate grading, or imposing an unstated minimal-context condition, would require reassessment of the source mapping. The read source does not impose minimality of the iid representation. Formal proof success alone does not settle these source or priority issues.

Independent final review and a root-owned source-only replay have passed. This is one scoped geometric-rate refutation; all supporting lemmas remain uncredited separately. The author label is "Research manuscript"; no worldwide priority is claimed.

References

[1] P. Cénac, B. Chauvin, C. Noûs, F. Paccaut and N. Pouyanne. *Variable Length Memory Chains: Characterization of stationary probability measures*. Bernoulli 27(3) (2021), 2011–2039. DOI: 10.3150/20-BEJ1299. Author-hosted published body: `https://pouyanne.perso.math.cnrs.fr/ccnppBernoulli.pdf`.

[2] Same authors and title. arXiv:2004.07893v1 (2020). `https://arxiv.org/abs/2004.07893v1`.

[3] Same authors. Supplement to [1], publication-hosted file `bej1299supp.pdf`, footer BEJ1299SUPP.tex, 11 November 2020; 20 pages. Public archive capture 20250506154122 of the publisher URL `https://projecteuclid.org/journals/supplementalcontent/10.3150/20-BEJ1299/bej1299supp.pdf`. SHA256 and acquisition limits are recorded in the packet provenance.

mathdb:2611.00026v1 04 Nov 2026