MATHAUDIT / DETAILED REPORT

Correctness, proof completeness, and novelty audit

Self-Referential Leading Digits of Exponential Sequences: Arithmetic Structure and Certified Search

Zihang Fang

Paper	arXiv:2607.23662v1
Source version	Submitted 26 July, 2026; announced July 2026.
Audit date	August 15, 2026
Prompt	Standardized prompt
Novelty cutoff	July 26, 2026
Paper license	arXiv non-exclusive distribution license
Scope	Statements, proofs, and novelty only

Important limitation. AI-generated analysis. This report is not a correctness certificate or a substitute for expert scrutiny. A Correct status may coexist with yellow typos or minor formal corrections that do not affect substantive soundness. It means that this audit found no unresolved substantive error under the stated criteria; it is not formal verification and should remain open to expert challenge and correction.

SUMMARY

01	**Statements**	Correct
02	**Proofs**	Correct
03	**Novelty**	No non-novelty findings

01 / STATEMENTS

Statements

Correct

The exact prefix and discrepancy identities, Lambert-layer geometry, subcritical distribution theorems, arithmetic classification and rigidity results, coherent-skeleton bounds, and certified sublinear locator have the scopes stated in the paper. The manuscript consistently keeps Lambert candidates and locator reports distinct from genuine prefix hits and does not claim to resolve the critical infinitude problem.

Theorems 3.1–3.8: Lambert layers, candidate count, and two-gap law

Pages 9–15 · Section 3 through Corollary 3.8 · arXiv:2607.23662v1

Correct

Inverting the two monotone endpoint functions gives exactly the half-open layer $[x_j, z_j)$ and its unique possible integer $m_j = \lceil x_j \rceil$. Lagrange inversion yields the convergent width expansion, while the integral representation of the inverse-root difference proves complete monotonicity. Since $x_{j+1} - x_j$ decreases to $1/\rho$ from above, the ceiling increments eventually take the two asserted values. Telescoping those increments and monotone inversion of Φ give the exact candidate counts, including the global 3–4 law and formula for $(c, b) = (2, 10)$. None of these steps identifies every candidate as an actual hit.

Verifiable source: Paper, version 1

Theorems 3.10 and 3.12: Deterministic subcritical distribution for algebraic parameters

Pages 15–19 · Lemma 3.9 and Theorems 3.10, 3.12 · arXiv:2607.23662v1

Correct

Matveev's lower bound gives an effective finite approximation type for $1/\rho$. The cited Tichy–Turnwald estimates provide ordinary discrepancy $O(N^{-\sigma})$ for every $\sigma < 1/(\eta + 1)$ and logarithmically weighted discrepancy $O(1/\log N)$. The exact Lambert root differs from $j/\rho + (\log j)/L$ by $O((\log j)/j)$; the dyadic endpoint-crossing argument preserves both estimates. Blocking the moving window at length $M^{1-(\sigma-\tau)/2}$ then balances its variation against discrepancy and gives the displayed main term and error for every $0 < \tau < \sigma$. The argument explicitly stops before the critical exponent $\tau = 1$.

Verifiable source: Tichy–Turnwald, logarithmic uniform distribution

Theorem 4.1 and Theorems 5.5–5.10: Exact counting identity and rational-slope classification

Pages 20–25 · Sections 4 and 5.1–5.2 · arXiv:2607.23662v1

Correct

The indicator is a difference of two adjacent floor values, and summation telescopes to the exact signed-discrepancy identity with all half-open endpoints preserved. For rational $\log_b c = P/Q$, substituting $n = Pm - Qk$ gives the exact residue parametrization. Valuation conditions characterize every integral residue branch, and for multiplicatively dependent integers the common-root representation $c = d^p$, $b = d^q$ reduces all hits to an eventually periodic congruence. The periodic map is balanced modulo q, yielding $A_{c,b}(X) = \log_b X + O(1)$.

Verifiable source: Paper, version 1

Theorems 5.12–5.30: Resonance shells, hit chains, and coherent skeletons

Pages 25–40 · Sections 5.3–5.5 · arXiv:2607.23662v1

Correct

Subtracting the phase relations for two hits places every fixed difference in one of the stated open shells, with strict inequalities matching the prefix convention. Irrationality-exponent bounds then give the asserted gap and representation estimates. At a positive-error floor center, direct rescaling proves nested prefix windows; within the exact nonempty range, the endpoint scales are forced to be coherent and every intermediate term follows. The Lambert formula for the coherent depth, its quadratic bound, continued-fraction lacunarity, and summability estimates correctly imply $T(X) = o(\sqrt{X})$ and the stated finite-type refinements. These are potential-skeleton bounds and do not assert that a center is a hit.

Verifiable source: Paper, version 1

Theorems 6.2–6.8: Certified continued-fraction locator and bit complexity

Pages 40–47 · Section 6 · arXiv:2607.23662v1

Correct

Adjacent convergents separate all block phases by $d_k = \|q_{k-1}\rho\|$. The outward rational-grid enclosure retains every hit and, under the half-margin condition, has total length below that separation, so at most one index is reported. Euclidean floor sums locate it without scanning. The interpolated envelope chooses the largest admissible length from the two neighboring convergents; finite irrationality type gives block length $\gg M^{1/\nu}$ and therefore $O(N^{1-1/\nu})$ blocks. Matveev separation and gcd-free equality testing justify polylogarithmic endpoint certification. Independent high-precision recomputation gives $d = 2.0276549588410307779\ldots \times 10^{-8}$ and threshold $\lceil q_k/(10^{d/2} - 1) - 1\rceil = 1{,}914{,}818{,}931{,}502{,}442$, agreeing with Corollary 6.8.

Verifiable source: Certified verifier repository

02 / PROOFS

Proofs

Correct

The proofs of the central statements and their material cited inputs are correct and complete. Exact endpoint conventions are propagated through the floor identities and resonance windows, the two distribution inputs are used within their proved finite-type range, and the search complexity includes certified transcendental comparisons and final exact verification.

Sections 2–4: Analytic inversion and discrepancy proof chain

Pages 8–20 · Propositions 2.1–3.3 and Theorems 3.5, 3.10, 3.12, 4.1 · arXiv:2607.23662v1

Correct and complete

The two real inverse branches are used only after both endpoint functions are increasing, and the half-open upper boundary is retained in the indicator and candidate criterion. The analytic implicit branch is identified uniquely before Lagrange inversion. The transfer from the model sequence to the Lambert roots counts all possible endpoint crossings uniformly on dyadic blocks, including atoms, and the moving-target summation balances both error sources with the claimed exponent. The signed-discrepancy formula is an exact telescoping identity rather than a probabilistic approximation.

Sections 5.1–5.4: Valuation, shell, and chain arguments

Pages 21–31 · Theorems 5.2, 5.5, 5.6, 5.10, 5.12–5.22 · arXiv:2607.23662v1

Correct and complete

The endpoint valuation elimination gives a genuine effective finite bound for independent integer parameters. In the dependent case, the eventual congruence period and induction give exactly one residue-class proportion $1/q$. For irrational slope, each double hit has a unique nonnegative shell index; the principal-shell criterion is applied only after its size hypothesis is verified. The chain-span estimate, reduction of $\lfloor q\alpha \rfloor/q$ to lowest terms, and Legendre criterion preserve the one-sided sign and denominator inequalities.

Section 5.5: Nested-window and coherent-depth arguments

Pages 32–40 · Theorems 5.23, 5.26, Proposition 5.28, and Theorem 5.30 · arXiv:2607.23662v1

Correct and complete

The floor identity $M = \lfloor q/(\varrho - 1) \rfloor$ orders every lower and upper endpoint exactly. For endpoint filling, comparing the two legal radix scales forces their integer difference to vanish; the strict coherent-range inequality then puts the common phase in every intermediate window. Convexity gives the unique positive Lambert root for $J_{\max}$, and the continued-fraction recurrence supplies both center lacunarity and the summable bounds used to derive total skeleton complexity.

Section 6: Packing, localization, and certification

Pages 40–47 · Theorems 6.2–6.6 and Corollaries 6.7–6.8 · arXiv:2607.23662v1

Correct and complete

All nonzero index differences inside a convergent block are strictly below the active denominator, so the stated best-approximation separation applies. The grid error budget includes phase rounding, outward endpoint rounding, and guard cells, leaving a strict reserve below d_k. Formula (152) is the exact residue-count identity used by the binary descent. The active-threshold proof handles both floor branches, while the bit model accounts for convergent construction, logarithmic sign tests, floor sums, reported-index verification, and output size.

Corollary 6.8: Reproducible numerical certificate

Pages 46–47 · Corollary 6.8 and Remark 6.9 · arXiv:2607.23662v1

Correct and complete

The cited public archive supplies a standalone Arb verifier, pinned top-level dependency, canonical JSON certificate, checksums, and byte-for-byte replay mode. The printed adjacent convergents have the required determinant and error direction, and direct high-precision recomputation places the real threshold strictly between 1,914,818,931,502,441 and 1,914,818,931,502,442. The corollary claims only a safe singleton report followed by exact verification, not existence of a hit.

03 / NOVELTY

Novelty

No non-novelty findings

No non-novelty findings.

METHOD

Scope and decision rules

This audit addressed only the correctness of central statements, the correctness and completeness of proofs, and evidence-backed non-novelty. It did not evaluate importance, exposition, style, suitability for publication, or author reputation.

An adverse finding is included only where the report can give a direct calculation, counterexample, formal contradiction, or a verifiable source establishing the precise claim. An inability to verify is kept distinct from a demonstrated defect. Technical compression is not treated as incompleteness when the omitted reasoning is recoverable and correct.

Before any finding was recorded, the argument was closed under immediate, mechanically checkable consequences of facts already established. An observation was suppressed when the printed conclusion followed by elementary logic, direct substitution, weakening, inclusion, or verified monotonicity and no substantive proof obligation or textual correction remained.

Typographical errors and minor formal corrections remain fully visible in yellow with their locations and repairs. When such a defect is unambiguous, local, and harmless to the paper's substantive conclusions, it does not lower the overall statements or proofs status from Correct.

The novelty panel reports only central results for which a verifiable primary source predating the earliest public version that contains the result proves the same statement up to notation, or from which the statement follows directly and immediately. It contains no positive novelty certifications and no entries based merely on an unsuccessful literature search.

All mathematical expressions in this report are compiled by LaTeX. The web presentation independently validates and renders the corresponding expressions with KaTeX.

Sources checked

1. Source paper: arXiv:2607.23662v1
2. Paper, version 1
3. Tichy–Turnwald, logarithmic uniform distribution
4. Certified verifier repository

arXiv acknowledgement

Thank you to arXiv for use of its open access interoperability. This service was not reviewed or approved by, nor does it necessarily express or reflect the policies or opinions of, arXiv.

APPENDIX

Standardized prompt

You are a mathematical auditor. Evaluate only the correctness of the paper's central statements, the correctness and completeness of its proofs, and evidence-backed non-novelty. Do not make editorial comments or evaluate importance, exposition, style, or suitability for publication. Keep the report neutral, precise, and professional.

PART 1 — CORRECTNESS OF THE MAIN STATEMENTS

Identify the paper's central statements and assess each one separately using exactly one of these outcomes:

- Correct: the statement has a correct, verifiable proof; or, despite a defect in the printed proof, you can give a fully verified repair or independent proof.
- Incorrect: you can give strong formal evidence that the statement as written is false, preferably a counterexample or a complete contradiction.
- Not able to verify: the supplied proof does not establish the statement, and you can neither supply a fully verified proof nor disprove it.

For every adverse or unverified outcome, restate the exact claim and hypotheses, identify its location in the reviewed version, and give the decisive reasoning. Do not infer that a theorem is false merely because its printed proof is defective.

Assign the overall statements status by the following priority:

1. Contains wrong statements, if at least one central statement is Incorrect.
2. Contains unsupported statements, if no central statement is Incorrect but at least one is Not able to verify.
3. Correct, if every central statement is Correct, apart from any typographical or minor formal findings allowed by Part 4.

PART 2 — CORRECTNESS AND COMPLETENESS OF THE PROOFS

Audit the proofs of the central statements and every material internal or cited input on which they depend. For each proof or proof segment, use the narrowest supported outcome:

- Correct and complete: every substantive step is verifiable and no case or crucial nontrivial step is missing.
- Incorrect as written: a displayed calculation, inference, lemma, or cited application is formally false or inapplicable.
- Incomplete as written: a necessary substantive step or case is missing, even though no explicit contradiction has been proved.
- Not able to verify: the available paper and sources are insufficient to decide whether the step is valid.

Do not flag omitted routine details, compressed arguments, or "analogously" steps when the omitted reasoning is recoverable and correct. For every defect, identify its smallest exact location, explain its downstream dependency, and classify any proposed repair as one of:

- Verified repair: supply enough reasoning to verify that the repair works throughout every affected use.
- Plausible repair only: describe the possible repair but do not treat it as established.
- No repair supplied.

Assign the overall proofs status by the following priority:

1. Contains incorrect or incomplete proofs, if at least one proof is Incorrect as written or Incomplete as written.
2. Contains unverified proofs, if there is no proved incorrectness or incompleteness but at least one material proof is Not able to verify.
3. Correct, if all audited proofs are Correct and complete, apart from any typographical or minor formal findings allowed by Part 4.

PART 3 — NOVELTY OF THE MAIN STATEMENTS

This is a one-sided, evidence-only check. Evaluate only central results. Report a result only when a verifiable primary source, publicly available before the earliest public version of the paper that contains the result, proves the same statement up to a change of notation, or when the result follows directly and immediately from such an earlier theorem.

For each non-novelty finding, provide the earlier source, its public date, the exact theorem or proposition, and an explicit hypothesis-and-conclusion mapping. A later source cannot establish non-novelty. Do not flag a result merely because an easier proof exists or because it follows through a nontrivial argument. Do not flag a result that the paper itself explicitly presents as known, a survey item, or an alternative proof.

Mention only statements that meet this non-novelty threshold. Say nothing about statements for which no such evidence was found: do not certify novelty, list unsuccessful searches, or add statement-level "no prior source found" entries. If the list is empty, the section may say only "No non-novelty findings."

PART 4 — TECHNICAL AND EVIDENTIARY RULES

1. Criticism must never be overstated. Call a statement incorrect, a proof incorrect or incomplete, or a result non-new only when strong formal evidence supports that exact judgment. Otherwise use Not able to verify where that outcome exists, or remain silent in the novelty section.

2. Review a specified paper version. Give exact page, theorem, equation, or proof locations; restate relevant quantifiers and hypotheses; and check domains, endpoints, parameter ranges, constants, and dependencies explicitly. Apply a defect only to the regimes and downstream results that actually depend on it.

3. Immediate-consequence closure: when deciding whether a conclusion is supported, treat the argument as including every immediate formal consequence of facts established earlier in the paper and available at that point. An immediate formal consequence is a local, mechanically checkable inference using only definitions already in force, elementary logic, direct substitution or renaming, weakening of an inequality or set inclusion, or a monotonicity or symmetry property explicit in a definition or already proved.

4. Such an inference must introduce no new lemma, estimate, construction, case analysis, limiting or compactness argument, external result, or new issue of uniformity or parameter dependence. Check the direction of monotonicity and the dependence of every threshold or constant explicitly. Do not exchange pointwise and uniform conclusions, change exceptional sets, interchange quantifiers or limits, or extend an endpoint under this rule unless the required fact was already established.

5. If the claimed conclusion follows after immediate-consequence closure, it is fully supported. Do not report a gap, unsupported statement, minor formal correction, or recommendation to add an explanatory sentence, and do not mention the suppressed item in the report. This suppression does not apply to a genuine typographical error as defined below, because a typo is a literal notation defect that still requires a textual correction. For example, if $P(delta)$ is proved for every sufficiently small delta > 0 and P is immediately monotone in the required direction, then the claim that $P(delta)$ holds for every delta > 0 requires no correction; the omitted monotonicity step is exposition rather than an audit finding.

6. Immediate-consequence closure permits reconstruction of an omitted inference; it does not permit silent alteration of the printed domain, hypotheses, or quantifiers of an intermediate statement later invoked outside its stated scope. If formal validity requires changing the text of that statement or its stated range, report the appropriate finding even when the repair is easy. Thus a result stated only for $t > 0$ and subsequently applied at $t = 0$ remains a formal scope mismatch unless the endpoint case was separately established before the application.

7. Apply this mandatory gate before recording any adverse, unverified, typographical, or minor formal finding: (a) identify the exact assertion or inference alleged to be defective; (b) reconstruct the strongest conclusion following from the preceding argument after all admissible immediate formal consequences are added; (c) identify the specific defect that remains — a formally demonstrated falsehood, a concrete nontrivial unresolved proof obligation, or a genuine domain, hypothesis, quantifier, or notation mismatch requiring a textual correction; and (d) state the evidence and minimal repair. If no specific defect remains, or if the only possible addition would be explanatory, suppress the finding entirely. Use Not able to verify only when a precise, nontrivial unresolved obligation can be identified.

8. A typographical error is an incorrect symbol, index, exponent, sign, variable name, or notation mismatch whose unique intended correction is mechanically determined from the surrounding text and does not alter the mathematical argument. Report it in yellow as Typo, with the printed text, correction, exact location, and reason it is harmless. It does not lower an overall status.

9. A minor formal correction is a boundary, endpoint, or displayed parameter-range defect that can be repaired consistently in every affected statement by a local change, without a new argument, without invalidating any related result or application, and without creating another substantive defect. Report it in yellow as Minor formal correction, with the same details. It does not lower an overall status.

10. Do not use either minor category if more than one materially different repair is possible; if the repair changes the substance or scope of a central result; if an excluded case is used later; if the correction fails to propagate consistently; or if a new mathematical idea is required. In such cases use the strongest substantive classification actually supported by evidence.

11. Write every mathematical expression as valid, compilable LaTeX. Use ... for inline mathematics and

...

for displayed mathematics, with balanced delimiters and braces. Do not leave raw LaTeX commands outside math delimiters.

12. Separate theorem truth from the correctness of the printed proof. A theorem can be Correct in Part 1 because of a fully verified independent proof while its printed proof is adverse in Part 2. Conversely, a defective proof alone never licenses an Incorrect theorem classification.

Organize the report under Parts 1–3 and apply Part 4 throughout. Do not add editorial commentary, scores, rankings, recommendations, or a list of suppressed observations.

END OF REPORT

Each claim in this prototype report remains open to independent challenge and revision.