MATHAUDIT / DETAILED REPORT

Correctness, proof completeness, and novelty audit

K-spherical horospherical averages on the Nagao quotient: tree combinatorics and exact discrepancy

Sanghoon Kwon

Paper	arXiv:2607.08704v1
Source version	Submitted 9 July, 2026; announced July 2026.
Audit date	August 15, 2026
Prompt	Standardized prompt
Novelty cutoff	July 9, 2026
Paper license	CC BY 4.0
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 central statements checked are correct. The exact rooted-shadow identities, the abstract spherical-shadow principle, the compact-orbit application, eventual exactness for compactly supported spherical tests, and the dense-truncation estimates all follow from the supplied combinatorics and standard Bruhat–Tits-tree coding with the stated normalizations.

Theorem 1.1: Exact rooted-shadow discrepancy

Pages 3 and 7–12 · Theorem 1.1 and Section 3 · arXiv:2607.08704v1

Correct

The shell count at height one normalizes to the probability law ρ^{ev}. Re-rooting a backward shadow gives its exact top-shell correction and missing tail, while partitioning a forward shadow by its first turn gives the displayed geometric mixture. These identities imply finite-window exactness, $L^1(\rho^{\mathrm{ev}})$ convergence for fixed roots, the three exponential regimes, and the uniform moving-root total-variation bound.

Verifiable source: Full paper, version 1

Theorem 1.2: Spherical-shadow principle

Pages 3–4 · Theorem 1.2 · arXiv:2607.08704v1

Correct

Testing a right K-invariant observable reduces the discrepancy exactly to its height profile against the pushed-forward measure. Convexity and the componentwise total-variation estimate from Proposition 3.15 yield $O(\|F\|_\infty q^{-M_T/2})$, and the backward formula yields the sharper $O(\|F\|_\infty q^{-M_T})$ rate. Proposition 2.3 identifies the limiting height law with the spherical pushforward of μ_X.

Verifiable source: Full paper, version 1

Theorem 1.3(i) and Theorems 4.4–4.5: Expanding translates of compact unipotent orbits

Pages 4 and 12–15 · Theorem 1.3(i), Lemma 4.2, and Theorems 4.4–4.5 · arXiv:2607.08704v1

Correct

Compactness restricts the initial spherical projection to finitely many sectors. In each sector, quotienting by $U \cap a_n K a_n^{-1}$ produces equal-mass cylinder cells forming a complete terminal descendant layer. The diagonal element translates by distance $2n$ in the type-preserving tree, while the finitely many initial roots contribute only bounded offsets, so the minimal cutoff tends to infinity. Theorem 1.2 gives equidistribution, and Proposition 3.11 makes the discrepancy exactly zero once the cutoff exceeds the finite support of the test profile.

Verifiable source: Full paper, version 1

Theorem 1.3(ii) and Theorems 5.7–5.10: Dense-orbit truncations and continued-fraction rates

Pages 4 and 15–21 · Theorem 1.3(ii) and Section 5 · arXiv:2607.08704v1

Correct

The lifted terminal layer consists of q^N equal-measure compact-open U-cylinders, so its spherical average is exactly a rooted-shadow distribution with multiplicity. Function-field continued fractions code alternating cusp excursions: an excursion of degree a_r has type-preserving length $2a_r$, and the turn contributes two further steps, giving $B_{r+1} = B_r + 2(a_r + 1)$. An irrational boundary point has infinitely many such blocks, hence $M_N \to \infty$; the bounded, admissible L^1, and effective estimates then follow from the corresponding rooted-shadow formulas.

Verifiable source: Paulin–Shapira continued-fraction coding

02 / PROOFS

Proofs

Correct

The proofs of the central results and their material geometric inputs are correct and complete.

Propositions 3.7–3.15: Rooted shell formulas and uniform discrepancy estimates

Pages 9–12 · Sections 3.3–3.6 · arXiv:2607.08704v1

Correct and complete

The shell formula has total mass one and yields ρ^{ev} on every fixed finite height window. The first-turn weights sum to $1 - q^{-N}$, with the remaining mass at the always-forward atom. Splitting turns at $M/2$ gives the stated total-variation estimate, and the L^1 conclusions follow from the summable ρ^{ev} tail and geometric first-turn weights.

Verifiable source: Full paper, version 1

Lemma 4.2 and Proposition 4.3: Finite-sector realization of compact horospheres

Pages 13–14 · Lemma 4.2 and Proposition 4.3 · arXiv:2607.08704v1

Correct and complete

The compact set YK has finite image in the discrete Nagao ray. After selecting lifts in the corresponding finitely many cusp sectors, the double quotient is partitioned into cosets of the compact-open stabilizer. Along a descendant edge these stabilizers have index q, so every family is the full terminal layer of its rooted sector with equal Haar mass. Translation by a_n adds $2n$ to the depth and changes each root height by only a sector-dependent bounded amount. This verifies both the finite convex decomposition and the uniform divergence of its minimal cutoff.

Verifiable source: Full paper, version 1

Lemma 5.3 and Propositions 5.5–5.6: Cylinder realization and exact cutoff coding

Pages 17–19 · Lemma 5.3 and Propositions 5.5–5.6 · arXiv:2607.08704v1

Correct and complete

Conjugating the boundary point to infinity reduces the cylinder claim to compact-open additive subgroups of $\mathbb{F}_q((t^{-1}))$; each child stabilizer has index q, which proves disjointness, equal mass, and the count q^N. The standard cross-section has first-return time twice the degree of the next partial quotient. Tracking the backward and forward halves of that return gives a constant cutoff followed by increments of two, with the additional two-step turn between blocks. Therefore the displayed block bases are exact and tend to infinity for every infinite continued fraction.

Verifiable source: Paulin–Shapira, cross-section and return-time coding

Theorems 4.4–4.5 and 5.7–5.10: Applications of the rooted-shadow estimates

Pages 14–15 and 19–21 · Sections 4.2 and 5.4–5.5 · arXiv:2607.08704v1

Correct and complete

Each application invokes a previously established shadow realization with the correct cutoff. Compact support is equivalent to finite support of the spherical height profile, so finite-window exactness applies. In the dense case, the top-shell and late-turn hypotheses in Definition 5.8 are exactly the two terms left by Proposition 3.10; the proof separately controls both and preserves the stated dependence on M_N.

Verifiable source: Full paper, version 1

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. Full paper, version 1
2. Paulin–Shapira, cross-section and return-time coding

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 $P(\text{delta})$ is proved for every sufficiently small delta > 0 and P is immediately monotone in the required direction, then the claim that $P(\text{delta})$ 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.