MATHAUDIT / DETAILED REPORT

Correctness, proof completeness, and novelty audit

Product sets in sets of returns and positivity of symmetric ergodic averages

Vitaly Bergelson, Saúl Rodríguez-Martín

Paper	arXiv:2608.02873v1
Source version	Submitted 3 August, 2026; announced August 2026.
Audit date	August 18, 2026
Prompt	Standardized prompt
Novelty cutoff	August 3, 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	Contains unsupported statements
02	Proofs	Contains incorrect or incomplete proofs
03	Novelty	No non-novelty findings

01 / STATEMENTS

Statements

Contains unsupported statements

The general product-set construction, its Cartesian and polynomial consequences, the semidirect-product and finite-field matrix results, and the principal Heisenberg counterexamples are supported. The quantitative symmetric-averaging bound for every finitely generated nilpotent group in Theorem 1.27 is not verified: its proof identifies the squares of one Følner sequence with another by a density assertion that fails for subgroups of the unitriangular group. The paper's qualitative nilpotent product-set conclusion remains correct through the independent polynomial-recurrence proof.

Theorems 3.1, 1.5, and 1.16–1.25: Product and Cartesian-product sets in measurable and combinatorial return sets

Pages 3–9 and 16–29 · Sections 2–4 · arXiv:2608.02873v1

Correct

The intersectivity lemma selects a set B of the claimed upper density for which every finite intersection of the associated measurable sets has positive measure. Applying it to

$$Y_g = \bigcap_i T_{\phi_i(g)} Y \cap T^{-1}_{\psi_i(g)} Y$$

gives $\psi_i(B)\phi_j(B) \subseteq R^T_\mu(Y)$ by one action identity. The amenable-group, Reiter-sequence, logarithmic-density, free-group, and convolution-power corollaries then follow from the appropriate mean ergodic theorem and the stated correspondence principle. The left/right order of every difference set is preserved in these deductions.

Theorems 1.46 and 1.48: Polynomial return sets and polynomial product-set consequences

Pages 14–15 and 30–39 · Section 5.1 and Section 5.2 · arXiv:2608.02873v1

Correct

Lemma 5.10 places the polynomial sequences generated by any IP-system inside a finitely generated VIP group. The nilpotent IP-polynomial recurrence theorem therefore makes the multiple-return set IP-star. Zorin-Kranich's norm-convergence theorem gives a Følner-sequence-independent limit; Lemma 5.14 rules out zero by extracting an IP-system from any zero-average sequence. Theorem 3.1 then converts this positive multiple average into all product-set inclusions in Theorem 1.46, and the Cartesian version follows in $H \times H$.

Verifiable source: Zorin-Kranich, nilpotent IP polynomial multiple recurrence

Theorem 1.27: Uniform quantitative positivity for finitely generated nilpotent groups

Pages 9 and 39–42 · Theorem 1.27 and its proof · arXiv:2608.02873v1

Not able to verify

The theorem asserts a system-independent bound

$$\lim_N \frac{1}{|F_N|} \sum_{g \in F_N} \mu(Y \cap T_{g^2} Y) \geq \lambda_G \mu(Y)^2$$

for every left or right Følner sequence, with the displayed explicit choice of λ_G. The proof's decisive approximation of a Følner average over B^G by squares of elements of A^G is false under the only established assumption that G is a subgroup of $\mathrm{UT}_n(\mathbb{Z})$; Part 2 gives a concrete subgroup where the alleged relative symmetric difference tends to 1/2. No counterexample to the averaging theorem itself is obtained, and the earlier polynomial-recurrence theorem proves positivity without this uniform constant, but the manuscript supplies no independent proof of the quantitative statement.

Theorems 1.31, 1.33, and 1.34: Symmetric averaging for semidirect products and matrix groups

Pages 10–11 and 43–47 · Section 6 · arXiv:2608.02873v1

Correct

For $H \rtimes K$ with H abelian, the square of kh has H-component $\phi_k(h)\phi_{k^2}(h)$, an endomorphism in h. Applying the mean ergodic inequality first in H and then the SAR hypothesis in K, followed by a diagonal Følner construction, gives the squared lower bound. Iteration handles $\mathrm{UT}_n(R)$. For $\mathrm{GL}_n(Q)$ over an algebraic extension of a finite field, the finite subgroups form a two-sided Følner sequence; regular split matrices occupy asymptotic proportion $1/n!$ and decompose into diagonal abelian subgroups, giving the claimed $\mu(Y)^2/n!$ bound.

Theorems 7.1 and 7.5: Heisenberg counterexamples separating left and right product phenomena

Pages 48–52 · Section 7 · arXiv:2608.02873v1

Correct

The rapidly separated Heisenberg boxes form a left Følner sequence and their pairwise quotients fall into three mutually recognizable coordinate regimes. Poincaré recurrence supplies a common nonzero central shift inside two positive-left-density sets; comparing the original and shifted products contradicts each of the three quotient regimes. The paired-box construction in Theorem 7.5 similarly forces the second middle coordinate to be bounded once the first is fixed, contradicting positive upper Banach density. These arguments establish the two advertised counterexamples.

02 / PROOFS

Proofs

Contains incorrect or incomplete proofs

The main intersectivity and polynomial-recurrence chains are correct. The proof of Theorem 1.27 contains a false asymptotic-density identity and therefore does not establish its quantitative nilpotent bound. A separate finite-set avoidance formula in Proposition 7.6 is also written in the wrong order, but it has a verified local repair and does not affect the principal Heisenberg counterexamples.

Proof of Theorem 1.27: Injectivity of squaring does not make its image asymptotically all of B^G

Page 41 · display beginning $|\operatorname{Sq}(A^G_{N_k})\triangle B^G_{N_k+1}|$ · arXiv:2608.02873v1

Incorrect as written

From $\operatorname{Sq}(A^G_N) \subseteq B^G_{N+1}$ and injectivity, the proof writes

$$\frac{|\operatorname{Sq}(A^G_N)\triangle B^G_{N+1}|}{|B^G_{N+1}|} = \frac{|B^G_{N+1}| - |A^G_N|}{|B^G_{N+1}|} \longrightarrow 0.$$

The preceding choice of N_k controls only the ratios $|A^G_{N_k+1}|/|A^G_{N_k}|$ and $|B^G_{N_k+1}|/|B^G_{N_k}|$; it gives no comparison between the A- and B-counts. For a concrete failure, take $n = 2$ and

$$G = \{I + 4mE_{12} : m \in \mathbb{Z}\} \le \mathrm{UT}_2(\mathbb{Z}).$$

Then A^G_N consists of the even choices of the parameter a, whereas every parameter b occurs in B^G_N, so $|A^G_N|/|B^G_{N+1}| \to 1/2$ and the displayed symmetric-difference ratio tends to 1/2. Downstream dependency: the comparison with the mean ergodic average on B^G, and hence the torsion-free case of Theorem 1.27 and its quotient step. Repair classification: No repair supplied. A repair would need a specially proved embedding or congruence subgroup on which the square image has full relative density, or a different averaging argument.

Theorem 1.48 and Theorem 1.46: The IP-star and positive-average argument closes the qualitative polynomial results

Pages 30–39 · Lemmas 5.10–5.14 and proofs of Theorems 1.48 and 1.46 · arXiv:2608.02873v1

Correct and complete

The derivative identities in Lemma 5.10 keep the generated family inside a finitely generated nilpotent polynomial group, and the separated-index convention used for the IP derivatives is checked at every material use of the cited recurrence theorem. Left and right Følner limits agree by inversion and a two-sided Følner sequence. If the common limit were zero, Lemma 5.14 constructs an IP-system on which the multiple-intersection values converge to zero, contradicting the positive IP-limit on a sub-IP-ring. This also independently proves the qualitative finitely generated nilpotent conclusion obtained from Theorem 1.46 with $p = q = \mathrm{id}$.

Proposition 7.6: The recursive forbidden set does not imply the claimed cross-disjointness

Page 52 · final paragraph of the proof of Proposition 7.6 · arXiv:2608.02873v1

Incorrect as written

To ensure $F_N g_N \cap F_M g_M a = \varnothing$ for every pair M, N, the proof says that it suffices to choose

$$g_N \notin E_N \cup \bigcup_{n<N} F_n g_n F_N^{-1}.$$

That set has the factors in the wrong order and omits the two translates by a needed for the two orientations of a cross-pair. Repair classification: Verified repair. At stage N, choose g_N outside

$$E_N \cup \bigcup_{n<N} F_N^{-1} F_n g_n \{a, a^{-1}\}.$$

The first set has zero upper Banach density and the second is finite, so such a choice exists; the two elementary rearrangements of a hypothetical intersection then give exactly the two forbidden translates. This proves $A \cap Aa = \varnothing$ and completes Proposition 7.6 without changing its statement.

Sections 2–4 and 6: Intersectivity, correspondence, and iterated averaging arguments

Pages 16–29 and 43–47 · Sections 2–4 and 6 · arXiv:2608.02873v1

Correct and complete

Fatou's lemma supplies the density-preserving point in the intersectivity lemma, and the null exceptional sets are removed simultaneously because the index set is countable. Every later product inclusion is a direct substitution into Theorem 3.1. The semidirect-product Følner sequence is diagonalized against finitely many translations at each stage; Jensen's inequality gives the second SAR exponent. The finite-field matrix count uses disjoint regular split eigenspace decompositions and loses exactly the factor $n!$. No additional unresolved hypothesis or case was found in these proof chains.

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:2608.02873v1
2. Zorin-Kranich, nilpotent IP polynomial multiple recurrence

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.