Formalization notes for Zhou's rigidity construction

Utensil Song

August 12, 2026

1 Introduction

1.1 Connes' rigidity conjecture

For a countable discrete group Γ, the group von Neumann algebra $L(\Gamma)$ is generated by the left regular representation of Γ. When Γ has **infinite conjugacy classes (ICC)**, meaning that every nonidentity conjugacy class is infinite, $L(\Gamma)$ is a II_1 factor. The reconstruction question asks how much of Γ remains visible in that analytic completion.

Kazhdan's property (T) says that a unitary representation with almost invariant unit vectors has a nonzero invariant vector [Ershov and Jaikin-Zapirain(2010), Section 3.1]. ICC makes the group von Neumann algebra a factor, while property (T) imposes representation-theoretic rigidity. These are the two hypotheses in the conjecture and in all three counterexample constructions below [Zhou(2026), Section 1, pp. 1–2].

Connes first proved that, for an ICC property-(T) group, $\mathrm{Out}(L(\Gamma))$ and the fundamental group $\mathcal{F}(L(\Gamma))$ are countable [Zhou(2026), Section 1, p. 1].

He subsequently proposed that such groups should be W^*-superrigid: for every countable group Λ, an isomorphism $L(\Lambda) \cong L(\Gamma)$ should force $\Lambda \cong \Gamma$. Zhou's introduction states this original 1982 conjecture and distinguishes it from the later, narrower rigidity question for higher-rank lattices [Zhou(2026), Section 1, pp. 1–2].

The counterexamples considered here disprove the broad conjecture. They do not settle that higher-rank-lattice question.

1.2 Counterexample background

The OpenAI and Zhou counterexamples were obtained independently and at nearly the same time. Zhou records that his preliminary manuscript predated the public OpenAI announcement, while acknowledging that the two projects were concurrent [Zhou(2026), Section 1, pp. 2–3].

OpenAI gives a countable family of pairwise nonisomorphic ICC property-(T) groups with isomorphic group factors [OpenAI(2026), Chapter 4]. Zhou gives an explicit pair through a different action-changing construction [Zhou(2026), Theorem A and Section 3].

The later Anthropic proof gives another explicit pair by comparing a split extension with a nonsplit one [Anthropic(2026), Sections 3–5].

These are not three presentations of one proof. The algebraic datum hidden by the factor sits at a different level in each construction, and each factor argument removes it by a different measurable mechanism. The mathematical and formalization motivations for comparing them are stated next. The detailed OpenAI and Anthropic architectures appear in § 9.1 and § 10.1; Sections 2–7 below follow Zhou's proof in its original order.

1.3 Motivation

The motivation for this formalization experiment is threefold.

1.3.1 Mathematical motivation

The three counterexamples form a controlled test of nonfaithfulness. OpenAI varies the abelian kernel [OpenAI(2026), Chapter 4]; Zhou fixes the kernel while varying the action and module structure [Zhou(2026), Sections 1 and 3]; and Anthropic fixes both while varying the extension class [Anthropic(2026), Sections 3–5]. In each construction, a genuine group distinction survives algebraically but disappears after measurable crossed-product transport. The comparison asks which distinctions survive representation, twisting, duality, and analytic completion. Property (T) here serves as a control condition: failure of reconstruction persists for groups with strong representation-theoretic rigidity.

1.3.2 Formalization motivation

Formalization asks where the information is lost. Named interfaces for the characteristic-two retraction, Fourier shear, measurable transport, continuous closure, spectral detection, and characteristic-module obstruction turn the claim that a factor forgets structure into an auditable dependency map. This separates direct translation from designed bridges and makes the project a consumer of Mathlib and TauCeti. Missing or bespoke interfaces mark candidates for reusable contributions to those libraries.

1.3.3 Agent-assisted work and human examination

The project also tests whether an agent-authored formalization can be brought to a standard suitable for human examination and digestion. Literature grounding, the exact external premise boundary, provenance, technical proof replay, mathematical audit, and companion notes answer different questions. Technical verification checks the formal statement and its replay; source and mathematical review ask whether the theorem, definitions, hypotheses, and supporting facts match the intended mathematics. A reader can then trace a claim from the literature to its Lean declaration, downstream use, audit status, and executable check.

1.4 Formalization design

The Lean development turns the preceding motivations into one construction stage and four certificate lanes: factor equivalence, property (T), ICC, and group nonisomorphism. The final theorem assembles those four endpoints.

The diagram is a map of both the proof and the notes. Solid arrows are internal Lean developments. The dashed arrow is the sole substantive mathematical input left external: the Ershov–Jaikin-Zapirain–Kassabov property-(T) theorem for $\mathrm{EL}_3(\mathbb{F}_2[t])$. Each lane uses a separate port at the construction and theorem nodes.

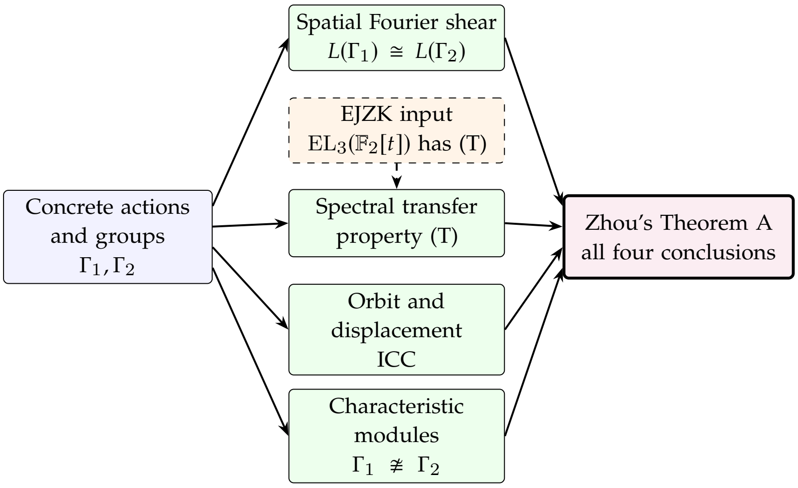

The mathematical weight is uneven:

- **Construction.** The characteristic-two calculation in Lemma 2.1.1, with its Lean realization in § 2.2, is concrete algebra and supplies every later carrier.
- **Factor equivalence.** Theorem 3.1.1, § 3.2, and § 3.3 carry heavy operator algebra: compact duals, Haar measure, Fourier transforms, von Neumann closures, and trace transport.
- **Property (T).** Definition 4.1.1, § 4.2, Theorem 4.3.1, § 4.4, Lemma 4.5.1, and § 4.6 carry heavy harmonic and representation theory. Spectral measures and finite detectors are internal; EJZK enters at one boundary.
- **ICC.** Theorem 5.1.1 is the lightest lane conceptually, but its implementation in § 5.2 still needs exact orbit and finite-quotient displacement calculations.
- **Nonisomorphism.** Theorem 6.1.1 is the deepest internal algebraic chain. Characteristic kernels, quotient twists, semisimplicity, and a finite cocycle obstruction must survive an arbitrary isomorphism.

Most definitions and algebraic laws translate directly. Six interfaces require special design:

- the square-span certificate implemented in § 2.2;
- the Fourier-shear proof of factor equivalence: its unitary implementation in Theorem 3.1.1, its automatic projection-order transport in § 8.3, the four conjugation identities in § 8.4, and the measure and closure arguments in § 3.2 and § 3.3;

- the quantitative detector and qualitative property-(T) boundary in Definition 4.1.1, § 4.2, Theorem 4.3.1, and § 4.4;
- the three-case ICC criterion in Theorem 5.1.1 and § 5.2;
- the quotient twist in § 6.2; and
- the proposition-valued assembly in § 7.1.

The formalization lessons are discussed in § 8.1, § 8.2, § 8.3, § 8.4, and § 8.5.

The code has two layers. Reusable group theory, linear algebra, topology, measure theory, and operator algebra live in a foundation layer. Zhou-specific carriers and certificates live in construction and paper layers. This separation lets the paper-facing modules state small consumer interfaces while coordinate or finite calculations remain behind named certificates.

The public formalization plan and statement map give the complete module and declaration crosswalk.

Each lane exports the smallest object needed by its next consumer: a spatial witness, a qualitative property-(T) certificate, orbit and displacement data, or a twist-stable module obstruction.

Four corrections were mathematically decisive:

- replace identity-action scaffolds by Zhou's actions;
- make the ICC abstraction handle the finite $\mathrm{Sp}_4(\mathbb{F}_2)$ factor;
- follow the quotient automorphism induced by a hypothetical group isomorphism; and
- tie the factor target to the Fourier shear and completed von Neumann closures.

The final interface remains proposition-valued. Its assembly and trust boundary are detailed in § 7.1 and Theorem 7.2.1; the remaining EJZK work is scoped in § 11.1, § 11.2, and § 11.3; the representation-universe boundary is recorded in § 11.4.

Section 1 separates the mathematical introduction, counterexample background, three motivations, and design. Sections 2–7 correspond respectively to Zhou's Sections 2–7; their headings record that correspondence without reusing Zhou's section number as the note title. Section 8 collects proof-engineering adaptations, Sections 9–10 compare the other two proof architectures, Section 11 scopes the remaining EJZK formalization, and Section 12 records the human–agent contribution and review workflow.

The formalization adapts selected public Lean proof blocks and organization patterns from a pinned `openai/ten-proofs` snapshot, but does not import that repository. The project records declaration-level adaptations in its public port map; the concrete Zhou carriers, corrected theorem boundaries, and final assembly are local to this development.

2 Characteristic-two tensor kernel (Zhou §2)

2.1 Tensor retraction in characteristic two

Lemma 2.1.1 (Characteristic-two tensor retraction `delta_diagonal` `concreteSquareSpanData`)
Let $k = \mathbb{F}_2$, $A = k[t]^3$, and let

$$C = (A \otimes_k A)^{\mathrm{Flip}}. \tag{2.1.1}$$

The diagonal function $\Delta(a) = a \otimes a$ lands in C. Although Δ is not a linear map, Zhou's characteristic-two calculation produces a surjective linear retraction $\delta : C \to A$ satisfying $\delta(\Delta(a)) = a$ [Zhou(2026), Section 2, Lemma 2.1]. The retraction is equivariant for the diagonal $\mathrm{SL}_3(k[t])$-action.

2.2 Formalizing the characteristic-two retraction

`delta_diagonal` `concreteSquareSpanData`

The Lean carriers and diagonal for Lemma 2.1.1 are `Connes.Construction.PaperKernel.C` and `Connes.Construction.PaperKernel.diagonal`; the coefficientwise map and retraction equation are `Connes.Construction.PaperKernel.delta` and `Connes.Construction.PaperKernel.delta_diagonal`.

The implementation fixes the standard polynomial and coordinate bases. It takes the coefficientwise Hadamard product on $k[t]$, applies it in the three coordinates of A, lifts the resulting bilinear map through the tensor product, and restricts it to C. The Boolean identity $x^2 = x$ gives $\delta(a \otimes a) = a$.

The basis-dependent formula implements a basis-free role: a linear left inverse to the diagonal squares, equivariant for the diagonal $\mathrm{SL}_3(k[t])$-action. Equivariance is not proved by expanding an arbitrary fixed tensor. The submodule `Connes.Construction.PaperKernel.squareSpan` is the span of the squares $a \otimes a$; the equation is immediate on each square, hence on their span.

The basis calculation in `Connes.Construction.PaperKernel.concreteSquareSpanData` proves that every flip-fixed tensor belongs to this span. The `SquareSpanData` interface therefore hides the coordinates from later sections. Consumers need only the spanning result, so a future basis-free construction could replace the formula without changing the action or theorem interfaces.

3 Factor equivalence by Fourier shear (Zhou §3)

3.1 Spatial factor equivalence

Theorem 3.1.1 (Spatial implementation of the factor equivalence `SpatialWitness` `isNormal` `tracialEquiv_of_spatialWitness` `paperGroupFactors_isomorphic`)
The group von Neumann algebra $L(G)$ is represented on $\ell^2(G)$ as the von Neumann closure of the left regular operators. Its canonical trace is the vacuum coefficient at δ_e. Consequently a unitary $U : \ell^2(G) \simeq \ell^2(H)$ gives the required tracial equivalence once two facts are proved: conjugation by U carries one closed operator algebra onto the other, and $U\delta_e = \delta_e$.

These are precisely the fields of `Connes.FactorWitness.SpatialWitness`.

The theorem `Connes.FactorWitness.tracialEquiv_of_spatialWitness` restricts unitary conjugation to the two group von Neumann algebras. `Connes.StarAlgEquiv.isNormal` records the automatic preservation of projection suprema by the induced star-algebra equivalence and its inverse. The equation $U\delta_e = \delta_e$ proves preservation of the canonical trace. Neither property is supplied as an extra assumption.

For Zhou's groups, the concrete unitary composes the two Fourier models with the fiber shear. The measure-transport and von Neumann closure obligations are separated in § 3.2 and § 3.3. Together they construct `Connes.PaperFactorClosure.paperSpatialUnitary`, `Connes.PaperFactorClosure.paperSpatialWitness`, and `Connes.PaperFactorClosure.paperGroupFactors_isomorphic`. This is the $*$-isomorphism asserted in [Zhou(2026), Proposition 3.4]; the Lean statement also records its projection-supremum preservation and preservation of the canonical trace.

3.2 Nonadditive shear and measure transport

`EquivariantHaarHomeomorph` `paperHaarHomeomorph`

Proposition 3.2 proves that the fiber shear preserves Haar measure and strictly conjugates the two actions. Its quadratic correction is generally not additive, so Zhou's Remark 3.3 does not claim a compact-group automorphism. Proposition 3.4 uses only the resulting equivariant measure-space isomorphism after Fourier transform [Zhou(2026), Proposition 3.2, Remark 3.3, and Proposition 3.4].

The independent crossed-product discussion in [OpenAI(2026), Chapter 4, Section 2.3, equation (2.1)] states the same abstract mechanism: group-law compatibility of the underlying compact spaces is not required.

The Lean boundary mirrors this distinction. The concrete `Connes.PaperCrossedHaar.paperFiberShearHomeomorph` and `Connes.PaperCrossedHaar.paperHaarHomeomorph` package continuity, measure preservation, and action equivariance, but no additive-homomorphism field.

At the foundation layer, `Connes.CrossedProduct.EquivariantHaarHomeomorph` projects this topological witness to the measurable transport interface. The extra topology is an implementation aid, not a stronger mathematical premise for the crossed-product identification.

A measurable equivariant isomorphism is mathematically sufficient for the abstract transport. The concrete homeomorphism is formally convenient because pullback preserves continuous coefficients. The proof uses that extra structure for the density bridge and forgets it at the measurable boundary, without falsely requiring an additive equivalence. The separate closure step is recorded in § 3.3.

3.3 Continuous coefficients and von Neumann closure

`continuousCrossedClosure_mem_iff`

Coordinate characters have norm-dense linear span in the continuous functions by Stone–Weierstrass. Since continuous coefficients act continuously as crossed-product multipliers, the generic continuous-coefficient closure theorem `Connes.CrossedProduct.vonNeumannClosure_crossedGeneratorSetOfContinuousCoefficients_eq` upgrades that density to equality of the generated von Neumann closures.

Pullback along the homeomorphism then gives `Connes.CrossedProduct.continuousCrossedClosure_mem_iff`. This does *not* say that continuous functions are norm-dense in L^∞; it compares the von Neumann closures generated after adjoining all continuous coefficients.

The paper-facing proof applies the transport and closure bridges once to each action, then composes the Fourier models with the shear. This removes parallel action-one/action-two closure plumbing while preserving Zhou's Proposition 3.4 endpoint and section order. Mapping a convenient generating family remains an intermediate step; the consumer-facing contract is equivalence of membership in the completed von Neumann algebras.

4 Property (T) by spectral transfer (Zhou §4)

4.1 Qualitative property (T)

Definition 4.1.1 (Qualitative property (T) interface `HasKazhdanPropertyT` `HasRelativePropertyT`)
Let G be a countable discrete group. A unitary representation π has **almost-invariant unit vectors** if every finite $K \subset G$ and every $\varepsilon > 0$ admit a unit vector ξ satisfying

$$\|\pi(g)\xi - \xi\| < \varepsilon \qquad (g \in K). \tag{4.1.1}$$

The project defines property (T) by requiring every such representation to contain a nonzero invariant vector. Relative property (T) for a subgroup $N \le G$ requires every such representation to contain a nonzero vector fixed by N.

4.2 Property-(T) transfer interface

`HasKazhdanPropertyT` `HasRelativePropertyT` `elementaryGroup`

The exact Lean boundary for Definition 4.1.1 uses the almost-invariance and property-(T) predicates `Connes.UnitaryRepresentation.HasAlmostInvariantUnitVectors`, `Connes.HasKazhdanPropertyT`, and `Connes.HasRelativePropertyT`. This qualitative form matches the final consumer, while quantitative Kazhdan sets and constants remain local to any future proof of the EJZK input.

The transfer argument uses the fixed subspace H^N of a normal subgroup N and its orthogonal complement. Relative property (T) produces a nonzero vector in H^N; property (T) of the quotient then acts on H^N. The relative-and-quotient transfer theorem `Connes.PropertyTTransfer.hasKazhdanPropertyT_of_relative_and_quotient` packages this step without exposing the concrete presentation of Zhou's groups to the spectral proof.

A future EJZK development should use Kazhdan pairs, ratios, or spectral gaps internally, then export `Connes.PaperPropertyT.elementaryGroup` for property (T) of the elementary group. Demanding a particular constant at the public boundary would strengthen the paper's hypothesis and couple every consumer to an arbitrary estimate.

4.3 Spectral criterion for split abelian extensions

Theorem 4.3.1 (Spectral bridge for split abelian extensions `spectral_criterion`)
A split extension $A \rtimes H$ with A abelian turns the restriction of a unitary representation to A into harmonic analysis on the Pontryagin dual $\widehat{A}$. For a vector ξ, a projection-valued spectral measure P yields the positive scalar measure

$$\mu_\xi(B) = \langle P(B)\xi, \xi\rangle. \tag{4.3.1}$$

Its total mass is $\|\xi\|^2$; kernel displacement becomes the energy $\int_{\widehat{A}} |\chi(a) - 1|^2 \, d\mu_\xi(\chi)$; and quotient-fixed vectors give invariant measures. A positive atom at the trivial character yields a nonzero invariant vector through its spectral projection.

4.4 Formalizing the spectral bridge `spectral_criterion`

The project records the group-theoretic data consumed by Theorem 4.3.1 in `Connes.SplitAbelianExtension`. The analytic layer can therefore use the kernel, quotient action, and section without unfolding a semidirect product.

The scalar measure and displacement interfaces are formalized in `Connes.ProjectionValuedSpectralMeasure`, and the finite-detection argument is packaged by `Connes.spectral_criterion`.

The direct consumer boundary is the following positive-atom estimate. If an invariant probability measure μ satisfies

$$c(1 - \mu(\{1\})) \le \sum_{a \in J} \int |\chi(a) - 1|^2 \, d\mu(\chi) \tag{4.4.1}$$

for some $c > 0$, then sufficiently small displacement on the finite detector J forces $\mu(\{1\}) > 0$.

Over $\mathbb{F}_2$, characters take values through the two-element circle subgroup. The local equivalence `Connes.BinaryPontryaginDual.pointwisePontryaginDualEquiv` identifies the dual of a linear dual with its evaluation module. This replaces implicit Fourier coordinates by an explicit additive equivalence; it is a formalization bridge, not an additional hypothesis.

4.5 Elementary generation of SL_3

Lemma 4.5.1 (Elementary matrices and special linear transport `elementarySubgroup_eq_top` `elementaryEquivSL3`)
Let $R = \mathbb{F}_2[t]$. The subgroup $E_3(R)$ is generated by transvections $I + rE_{ij}$ for $i \neq j$. Zhou's Proposition 4.1 first uses Euclidean row reduction to identify it with $SL_3(R)$, then invokes the property-(T) theorem of [Ershov et al.(2017), Theorem 1.1 and Section 1.2]. The equality step is not part of the remaining external boundary.

Consequently the multiplicative equivalence `Connes.PaperPropertyT.elementaryEquivSL3` transports property (T) from the elementary group named by the cited theorem to $SL_3(R)$.

4.6 Formalizing elementary generation and transport

`elementarySubgroup_eq_top` `elementaryEquivSL3`

The formalization proves the equality used in Lemma 4.5.1 by determinant-preserving row operations and Euclidean descent on polynomial entries, culminating in `Connes.SpecialLinear.elementarySubgroup_eq_top`.

It then constructs `Connes.PaperPropertyT.elementaryEquivSL3` and transports property (T) across the named multiplicative equivalence. The external premise therefore names the exact group in the cited theorem, while the paper-facing consumer works with $SL_3(R)$.

Rewriting subgroup equality into carrier equality would expose dependent coercions throughout Section 4. Keeping transport at one explicit equivalence also lets later finite-index and quotient arguments remain independent of the chosen presentation.

5 ICC by orbit displacement (Zhou §5)

5.1 ICC criterion

Theorem 5.1.1 (Orbit and displacement criterion for ICC `isICC_of_product_quotient` `paper_gammaOne_icc` `paper_gammaTwo_icc`)

For a countable group, the ICC condition asks that every nonidentity conjugacy class be infinite. In a semidirect product $N \rtimes (S \times Q)$, the proof can be separated from the concrete formulas by supplying two kinds of certificates:

1. every nonidentity $a \in N$ has an infinite S-orbit; and
2. every nonidentity $q \in Q$ admits $a \in N$ whose displacement $a(q \cdot a)^{-1}$ is nonidentity and has an infinite S-orbit.

The structure `Connes.PaperICC.ActionData` is exactly this direct-consumer boundary. The generic theorem `Connes.PaperICC.isICC_of_product_quotient` then proves that these certificates imply ICC by splitting a nonidentity element into three exhaustive cases.

- A nontrivial S-coordinate inherits an infinite conjugacy class from S.
- A pure kernel element uses the first certificate.
- A surviving Q-coordinate reduces to the second certificate.

This organization follows the cases used in [Zhou(2026), Section 5] without building the concrete tensors into the reusable ICC theorem.

5.2 Formalizing the ICC criterion

`isICC_of_product_quotient` `paper_gammaOne_icc` `paper_gammaTwo_icc`

The paper-specific work behind Theorem 5.1.1 is confined to orbit calculations in `ICCOrbits.lean`. Those calculations construct the action data and yield `Connes.PaperICC.paper_gammaOne_icc` and `Connes.PaperICC.paper_gammaTwo_icc`.

The displacement certificate is the finite-quotient input. Conjugating an element with nontrivial Q-coordinate by a kernel element changes its kernel coordinate by $a(q \cdot a)^{-1}$. Asking directly for an infinite family of conjugates would duplicate this algebra for each action. The certificate instead exposes the one displacement whose infinite S-orbit suffices.

6 Nonisomorphism by characteristic modules (Zhou §6)

6.1 Characteristic-module obstruction

Theorem 6.1.1 (Characteristic modules detect nonisomorphism `firstProduct_semisimple` `paperCharacteristicModuleEquiv` `paperGroups_not_isomorphic`)
The Section 6 invariant is semisimplicity of a module obtained from the characteristic abelian kernel. This is suitable for formal transport: a group isomorphism identifies the characteristic kernels, descends to an automorphism σ of the finite quotient, and induces a linear equivalence between the first quotient module and the second module with its action restricted along σ. Semisimplicity is invariant under that linear equivalence.

The characteristic-kernel and quotient-twist boundaries used here are separated in § 6.2.

On the first side, explicit decomposition into simple summands culminates in `Connes.PaperModuleSemisimple.firstProduct_semisimple`.

On the second side, a nonsplit submodule obstruction is proved for every quotient twist σ. The bridge `Connes.PaperNonisomorphism.paperCharacteristicModuleEquiv` turns a hypothetical group isomorphism into the forbidden module equivalence.

The public contradiction is `Connes.PaperModuleSemisimpleTransport.paperGroups_not_isomorphic`. This records the invariant used in [Zhou(2026), Section 6] while keeping the group-theoretic, finite-certificate, and module-theoretic arguments in separate files.

6.2 Characteristic kernels and quotient-twisted transport

The characteristic-kernel step maps an abelian normal subgroup to both factors of $\mathrm{SL}_3(\mathbb{F}_2[t]) \times \mathrm{Sp}_4(\mathbb{F}_2)$. It therefore needs two vanishing results:

- The polynomial-matrix argument `paperSL3NoNontrivialAbelianNormalSubgroup` for the polynomial-matrix factor handles the SL_3 image.
- The stronger finite result `Sp4.no_nontrivial_normal_abelian_subgroup` handles the symplectic image through a kernel-checked certificate.

The product argument applying both inputs is `paperSemidirectNormalAbelian_le_kernel`. This confines finite enumeration to the symplectic factor and keeps both factor-specific proofs out of the abstract module transport.

Design adaptation: twisting does not repair nonsplitting. Restriction along a group automorphism is an equivalence of module categories, with inverse given by restriction along the inverse automorphism. The formal proof nevertheless exports the exact twisted obstruction needed by the consumer, so no unproved categorical transport principle remains hidden at the final contradiction.

7 Theorem assembly and trust boundary (Zhou §7)

7.1 Exact external boundary

`theoremA` `elementaryGroup`

The public theorem is `Connes.theoremA`. Its sole mathematical parameter is $\mathrm{HasKazhdanPropertyT}(\mathrm{EL}_3(\mathbb{F}_2[t]))$. The corresponding Lean boundary consists of `Connes.PaperPropertyT.elementaryGroup` and `EJZKInput`.

All spectral, factor, ICC, and nonisomorphism certificates downstream of that parameter are constructed in the project. This is completeness for the load-bearing content of Zhou's Theorem A under the project's definitions, not a claim that every expository sentence of [Zhou(2026)] has been encoded.

The dependency boundary is deliberately narrow:

$$\boxed{\begin{array}{c}\mathrm{EL}_3(\mathbb{F}_2[t])\\ \text{has property (T)}\end{array}} \xrightarrow{\text{transport}} \boxed{\begin{array}{c}\mathrm{SL}_3(\mathbb{F}_2[t])\\ \text{has property (T)}\end{array}} \xrightarrow{\text{certificates}} \boxed{\begin{array}{c}\Gamma_1, \Gamma_2\\ \text{satisfy Theorem A}\end{array}}$$

The solution import graph contains no axiom, opaque declaration, sorry, or admit. Lean's axiom report for `Connes.theoremA` contains only:

- `propext;`

- `Classical.choice`; and
- `Quot.sound`.

These are the usual Lean foundations used by the imported library. The one `sorry` in an independent Comparator challenge is outside the theorem's import graph.

7.2 Theorem A

Theorem 7.2.1 (Assembly of Theorem A `theoremA` `theoremA`)
The final assembly has four independent certificate branches. The spectral branch proves property (T) for both concrete groups from the EJZK input and the internally constructed finite detectors. The other branches provide ICC, the trace-preserving group-factor equivalence, and nonisomorphism. The theorem `Connes.PaperTheoremACompletion.theoremA` only gathers those endpoints; it performs no additional mathematical construction.

The source-facing declaration `Connes.theoremA` unwraps the one-field EJZK input and states the full existence conclusion. Its single substantive mathematical premise is $T(\mathrm{EL}_3(\mathbb{F}_2[t]))$. The local equality $\mathrm{EL}_3 = \mathrm{SL}_3$, all transfers, and every certificate for the two groups are already proofs in the import graph.

Kernel and source trust boundary. The precise external premise, axiom report, and meaning of completeness are recorded once in § 7.1.

8 Formalization adaptations

8.1 Dependency placement and consumer boundaries

`completion_of_spectralData` `paperSpatialWitness` `ActionData`

Place instances at their mathematical source. The compact-dual topology, second-countability, and measurable-space instances belong beside the Pontryagin dual construction in Section 3, not at their first measure-theoretic use in Section 4. Likewise, the additive group instance for a representation module belongs at the module boundary in Section 6. Early placement makes Zhou's section order an acyclic dependency discipline instead of a directory convention.

Concrete examples are:

- the dual topology instances `paperKernelTopology` and `paperKernelDiscreteTopology`;
- the Borel instances `paperCharacterMeasurableSpace` and `paperCharacterBorelSpace`; and
- the representation-module carrier instance `representationAsModuleAddCommGroup`.

Isolate certificates from structural theorems. The finite symplectic enumeration, concrete orbit witnesses, and detector estimates live in dedicated files. Their consumers see named propositions or small structures. This prevents an exhaustive certificate, coordinate calculation, or normalization choice from contaminating a generic transfer theorem.

State the direct consumer boundary. The external theorem is requested as property (T) of the elementary group, not as a stronger numerical constant. The factor bridge asks for a spatial unitary that maps completed algebras and the vacuum. The ICC bridge asks for orbit and displacement data. Each boundary contains exactly what its next theorem consumes, which makes hypothesis strength and reuse auditable.

8.2 Structure forgetting and equivalence reuse

`EquivariantHaarHomeomorph` `continuousCrossedClosure_mem_iff`

The formalization uses the topology of the fiber shear at an intermediate step. It defines the shear as an equivariant measure-preserving homeomorphism because pullback must preserve continuous coefficients. Once this density argument is complete, the crossed-product theorem uses only the underlying equivariant measurable equivalence.

Zhou's shear need not be additive, so the Lean proof does not treat it as a compact-group automorphism. Continuity is retained for the density argument and then forgotten. This matches the two roles separated in [Zhou(2026), Proposition 3.2 and Remark 3.3].

A second lesson came from the reverse inclusion between the generated operator algebras. Mathlib's continuous pullback is already a star-algebra equivalence, so its surjectivity proves that inclusion without a separate involutivity calculation. Similarly, the inverse law for conjugation by a linear isometry equivalence replaces a pointwise operator calculation. In both cases, an equality that first looked computational follows from the inverse laws of an equivalence.

8.3 Normality of the factor isomorphism

`IsProjectionSupremum` `isNormal`

Zhou's Proposition 3.4 identifies the two group factors by combining Fourier transform with the fiber shear [Zhou(2026), Proposition 3.4]. The Lean development records two properties of the resulting star-algebra equivalence explicitly. Preservation of projection suprema in both directions follows automatically from the star-algebra equivalence between these concrete operator subalgebras, while preservation of the canonical trace is proved from the spatial unitary.

The definition `Connes.IsProjectionSupremum` says that p is a projection, every member of S is a projection below p, and p lies below every projection that is an upper bound for S. The last condition ranges only over projections, rather than all upper bounds as in the stronger ambient predicate `IsLUB`.

For star subalgebras of bounded operators, the usual order on projections is characterized by $p \le q \iff pq = p$. A star-algebra equivalence preserves multiplication, so `Connes.IsProjectionSupremum.map_starAlgEquiv` uses this characterization to preserve projection suprema. Applying the same argument to the equivalence and its inverse gives `Connes.StarAlgEquiv.isNormal`.

This refactor leaves the definitions of projection supremum and normal star-algebra equivalence independent of bounded operators. The operator order appears only in the theorem that applies those definitions to the group factors in Theorem 3.1.1.

8.4 Four conjugation identities

To identify the generated von Neumann algebras, Lean checks what each chosen unitary does to the regular-representation generators. There is one calculation for the abelian kernel and one for the acting group in each of Zhou's two group constructions, giving four identities in all.

For the first group, see the kernel `calculation` and the acting-group `calculation`. For the second group, see the corresponding kernel `calculation` and acting-group `calculation`.

These calculations spell out a step summarized by the Fourier and crossed-product identifications in [Zhou(2026), Proposition 3.4]. They carry the generators through the Fubini equivalence for a semidirect product and then through the Fourier unitary.

Most explicit heartbeat settings disappeared during the later proof cleanup. These four declarations still require a larger elaboration limit because Lean must simplify nested equivalences and pointwise operator formulas. The limit controls elaboration time; it adds no mathematical assumption.

8.5 Proposition-valued final assembly

`theoremA` `theoremA`

Keep final assembly proposition-valued. Section 7 combines previously proved endpoints. It does not accept a broad record containing the desired conclusion as a field. This "certificate isolation" makes `#print axioms` meaningful and lets semantic review trace each conjunct to the file where its mathematics occurs.

The nonisomorphism conjunct is stated directly as the absence of a multiplicative equivalence, $\neg\,\mathrm{Nonempty}(\Gamma_1 \simeq \Gamma_2)$, rather than through a project-local synonym. This is Lean's standard expression for the absence of a group isomorphism, so the result proved in Section 6 can be used directly.

The same assembly pattern applies when a linear paper proof crosses topology, analysis, finite computation, and algebraic transport: export small propositions from those layers, then combine them only at the theorem boundary.

9 OpenAI proof architecture

9.1 Changing the kernel while preserving the measured action

`paperHaarEquiv` `paperGroupFactors_isomorphic`

Chapter 4 of [OpenAI(2026)] constructs a countable family $\Gamma_n = E_n \rtimes K$. The acting property-(T) group K is fixed, while the discrete abelian kernels E_n are dual to different compact group laws on one underlying probability space. Binary carry changes the compact group structure but not its Borel probability space or the relevant K-action.

Pontryagin duality therefore presents every group factor by the same measured crossed product $L^\infty(X \times Y) \rtimes K$ [OpenAI(2026), Chapter 4, Sections 2.3, 4, and 6.1–6.2].

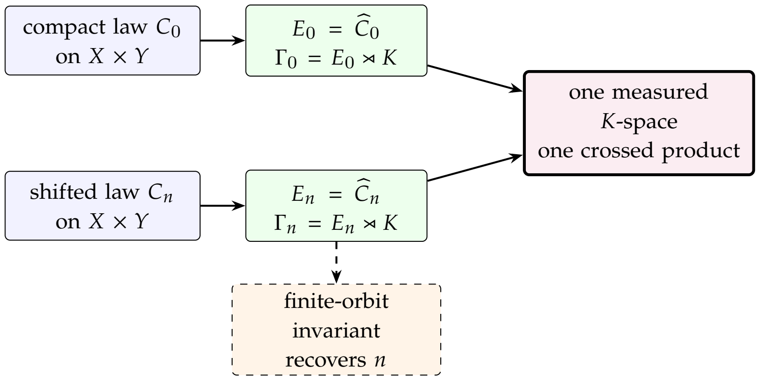

The proof separates the information forgotten by the group-factor functor from the information retained by the abstract group. The common measured action gives isomorphic factors.

For the first pair, order-four torsion already distinguishes the groups. For the full family, the parameter n is recovered from an intrinsic finite-orbit subgroup of $E_n[2]/2E_n$; finite-index embeddings simultaneously prove mutual commensurability [OpenAI(2026), Chapter 4, Sections 5.1, 6.2, and 6.4–6.5].

The same chapter proves property (T) by a relative-property-(T) spectral estimate and a Boolean-polynomial support bound, rather than taking it only from a classical lattice theorem [OpenAI(2026), Chapter 4, Sections 5.2–5.5].

This is not Zhou's deformation. Zhou keeps both the discrete kernel and the quotient fixed and changes the action; the quadratic shear then conjugates the two dual actions [Zhou(2026), Section 1, pp. 2–3]. Here the action on the underlying measured space is already common, while the compact group law, hence its discrete dual, changes.

The present Lean development formalizes Zhou's action-changing route. Its Fourier and crossed-product interfaces are conceptually reusable, but the shifted carry groups, the infinite family, and its finite-orbit invariant are not claims of the current formalization. This change of kernel is also what lets the OpenAI construction produce an infinite fiber, whereas Zhou's theorem selects one explicit pair.

The balance of proof weight is correspondingly different. Once the compact models are identified as the same measured K-space, the factor comparison is short. More of the work lies in constructing the shifted family, proving the initial relative-property-(T) estimate, and extracting an intrinsic invariant that distinguishes every parameter. Zhou instead puts substantial weight into the explicit action conjugacy and into transporting a module obstruction through an arbitrary group isomorphism.

10 Anthropic proof architecture

10.1 Changing the extension class and measurably gauging it away

`paperFiberShear_conjugates_paperActions` `paperGroupFactors_isomorphic`

The construction in [Anthropic(2026)] fixes both the kernel $A = \mathbb{Z}^4$ and quotient $Q = \mathrm{Sp}_4(\mathbb{Z})$. It compares the split semidirect product $\Gamma_0 = A \rtimes Q$ with an extension Γ_τ defined by a nonzero two-torsion class $[\tau] \in H^2(Q, A)$.

The algebraic difference lies neither in the kernel nor in the action, but in the extension cocycle. The class is obtained as a Bockstein of the theta-characteristic cocycle [Anthropic(2026), Sections 3.1–3.3].

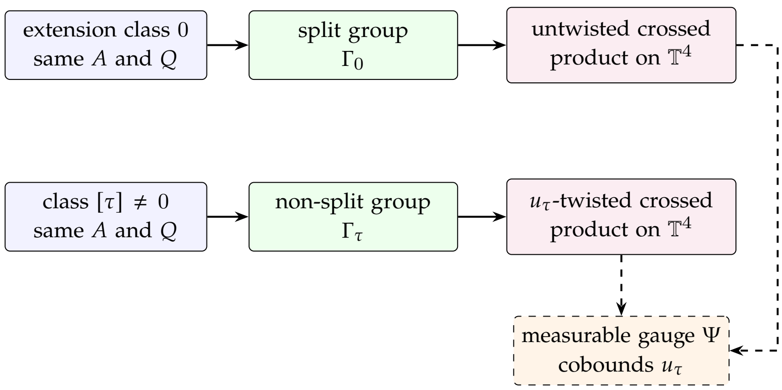

After Fourier transform, the two group factors become an untwisted and a twisted crossed product over $\mathbb{T}^4$. A Borel choice of representatives for $\mathbb{R}^4/\mathbb{Z}^4$ produces an integral carry cocycle. The paper combines that carry with a quadratic refinement of the symplectic form to construct phases Ψ_g; their coboundary is exactly the dual twist u_τ.

The resulting gauge is a trace-preserving isomorphism of the crossed products [Anthropic(2026), Sections 2.3 and 5.1–5.3, especially Theorems 5.1 and 5.2].

Nonisomorphism survives because $[\tau] \neq 0$: one extension splits and the other does not [Anthropic(2026), Section 4]. The paper proves ICC directly and obtains property (T) from the classical $\mathrm{Sp}_4(\mathbb{Z})$ and relative-property-(T) input, followed by finite-index permanence [Anthropic(2026), Lemmas 3.7–3.8].

The constructions vary three different algebraic data:

- OpenAI changes the kernel through a different compact dual law;

- Zhou changes the action through a nonsplit module extension; and
- Anthropic changes the group-extension class in H^2.

Their measurable repairs also differ: an already common measured action, an explicit conjugacy of actions, and a one-cochain that untwists a crossed-product two-cocycle. Remark 7.1 of [Anthropic(2026)] describes the last point as a torsion class becoming measurably trivial on the connected dual while its branch defects are repaired by the quadratic carry terms.

Formalizing this proof would not be a small refactor of the Zhou development. The current factor-transport layer treats untwisted crossed products and strict action conjugacy. The Anthropic route would need:

- explicit group-extension cohomology;
- cocycle-twisted crossed products;
- Borel fundamental-domain representatives; and
- a verified measurable gauge.

Its property-(T) lane appears shorter than the current EJZK boundary, but its factor lane introduces new infrastructure. These claims are not included in the present Lean theorem.

11 EJZK formalization outlook

11.1 EJZK route and effort estimate

`SimpleGraph.lapMatrix` `elementaryGroup`

The exact property-(T) input is classical and mathematically well established. The literature survey in [Ershov et al.(2017), Section 1.2, pp. 3–4] recalls that Kazhdan's 1967 theorem, together with subsequent work of Vaserstein, gives property (T) for $G_\Phi(\mathbb{F}[t]) = E_\Phi(\mathbb{F}[t])$ when $\mathbb{F}$ is finite and Φ is a reduced irreducible classical root system of rank at least two.

Taking $\Phi = A_2$ and $\mathbb{F} = \mathbb{F}_2$ yields the required $E_{A_2}(\mathbb{F}_2[t]) = \mathrm{EL}_3(\mathbb{F}_2[t])$ instance. The open boundary is therefore one of formalization completeness, not mathematical uncertainty.

The only remaining substantive premise is property (T) for $\mathrm{EL}_3(\mathbb{F}_2[t])$. The route estimated below specializes the 2010 elementary-group theorem to $R = \mathbb{F}_2[t]$

and $n = 3$. Its graph and six-root estimates are cited in the implementation layers below. The later root-graded treatment supplies broader supporting language.

Our central estimate remains **24–28 thousand new Lean lines**; a wider plausible range is **18–35 thousand**. That range is about 0.8–1.5 times the roughly 23,000 lines already in the project. We expect the proof search and engineering to require **two to four times the effort of the present formalization**.

This estimates one explicit, quantitative route that fits the present Lean interfaces; it is neither a measure of confidence in the theorem nor a lower bound for formalizing the classical route. The literature supports the route and its decomposition, not these LOC or effort figures.

The six implementation layers, with overlapping uncertainty ranges, are:

1. quantitative Kazhdan subsets, constants, and relative ratios, with a bridge to the current qualitative predicate; our internal allocation is **2–3.5 thousand lines**;
2. the six additive root-subgroup embeddings and the direct A_2 commutator relations; our internal allocation is **2–4 thousand**. A Steinberg-group detour is optional, not required for this direct matrix-group specialization;
3. fixed-space angles, codistance, and the class-two nilpotent estimate; our internal allocation is **5–9 thousand**;
4. the unweighted standard-Laplacian boundary case on six vertices. This is Section 5.4 of [Ershov and Jaikin-Zapirain(2010)]. Our internal allocation is **1.5–3 thousand**; the weighted criterion in Section 5.5 is a separate generalization;
5. the packaged relative-ratio estimate for the union of all six root subgroups in [Ershov and Jaikin-Zapirain(2010), Proposition 6.1, p. 34]; our internal allocation is **3–5.5 thousand**. Six repeated Lean instantiations would be an implementation choice, not six mathematical inputs; and
6. assembly plus the exact qualitative export theorem for the elementary-group property-(T) boundary; our internal allocation is **1.5–3 thousand**.

The allocations overlap. Their subtotal also excludes integration adapters, proof-search retries, and pin-specific repair risk, so it is not an arithmetic derivation of the 18–35 thousand-line envelope.

The existing analytic shell is not part of the missing estimate. `Connes.spectral_criterion_unconditional` already constructs the positive spectral functional and closes the scalar/PVM route from finite spectral detection to property (T).

The remaining work is the quantitative root-subgroup, fixed-space, graph, and detector mathematics listed above; it should reuse that shell rather than construct another PVM layer.

11.2 Available Mathlib and TauCeti infrastructure

`SimpleGraph.lapMatrix` `SimpleGraph.posSemidef_lapMatrix`

The pinned Connes project uses Mathlib revision `062f1e3d` with Lean 4.34.0-rc1. Mathlib supplies the graph Laplacian and its positivity theorem, exposed by the declaration links above, but not the Ershov–Jaikin-Zapirain graph-of-groups criterion, Kazhdan ratios, root gradings, codistance, or the graph argument.

Kassabov's estimates for a free associative integer ring require quotient transport before application to $\mathbb{F}_2[t]$ [Kassabov(2005), Theorem 2.2 and Lemmas 2.3–2.7]. For an arbitrary finitely generated ring and module, the cleaner calibration is [Ershov et al.(2017), Appendix Theorem A.1]. The direct route is the earlier quantitative development [Ershov and Jaikin-Zapirain(2010), Sections 2–6 and Appendix A].

TauCeti at revision `cc6ce75` uses Mathlib `de5ce8a` and Lean 4.34.0-rc1. Its diagonal-dominance proof consumes the same Mathlib Laplacian facts.

Its source and roadmap contain no direct property-(T), Kazhdan, codistance, root-graded, or magic-graph lane. It is therefore not a drop-in dependency for this boundary; aligning project pins would not reduce the missing mathematics.

11.3 Alternative EJZK routes and literature anchors

`elementaryGroup`

Bounded elementary generation is not presently shorter. Nica's [Nica(2019), Theorem 2, p. 2] gives

$$\nu_3 = \frac{3 \cdot 3^2 - 3}{2} + 29 = 41, \tag{11.3.1}$$

and its proof invokes the Kornblum–Artin theorem for irreducibles in polynomial arithmetic progressions, stated as [Nica(2019), Theorem 3, p. 2]. That input is absent from the pinned Mathlib revision.

Our internal estimate for a self-contained Nica–Kassabov route is **20–45 thousand lines** and **three to six times the current engineering effort**, with roughly 70 percent internal uncertainty.

Assuming Kornblum–Artin gives our internal estimate of **8–15 thousand lines**, but merely moves the external boundary. Our internal estimate for a reusable formalization of the generic EJK theorem is **45–90 thousand lines**; that generality is unnecessary for Zhou's application.

The route decomposition has four literature anchors:

- Zhou's exact invocation [Zhou(2026), Proposition 4.1, pp. 8–10];
- the elementary-group theorem, unweighted boundary graph, and packaged six-root estimate [Ershov and Jaikin-Zapirain(2010), Theorem 1.1, p. 1; Section 5.4, pp. 25–29; Proposition 6.1, pp. 34–35];
- the broader root-graded theorem and arbitrary-ring relative estimate [Ershov et al.(2017), Theorem 1.1, p. 3; Appendix Theorem A.1, p. 110]; and
- Kassabov's free-associative-ring estimates and Nica's exact-ring alternative.

None states our LOC or effort ranges.

11.4 Representation-space universe boundary

`HasKazhdanPropertyT`

The cited literature formulates property (T) by quantifying over unitary representations. The basic definitions and elementary-group theorem appear in [Ershov and Jaikin-Zapirain(2010), Section 3.1 and Theorem 1.1]; the root-graded generalization is [Ershov et al.(2017), Theorem 1.1]. Lean must also make the size of each carrier explicit.

The current `Connes.HasKazhdanPropertyT` takes a group in `Type u` and quantifies over Hilbert carriers in the same `Type u`. This conventional interface is sufficient for the theorem schema formalized here.

It matches the spectral consumer: `spectral_criterion_unconditional` is uniform in each `K : Type u` and each `UnitaryRepresentation G K`, rather than selecting one concrete ℓ^2 or regular representation.

The interface does not assert universe independence. A classical theorem quantifying over carriers in every universe implies this predicate by restriction to `Type u`; thus EJZK safely supplies the premise used here. The converse is nontrivial and is not part of the present interface.

For a countable group, that converse has a two-stage route:

1. Choose witnesses for every finite subset and a cofinal sequence of positive accuracies. The closed span of their group orbits is invariant, complete, separable, and still has almost-invariant unit vectors. An invariant vector there remains invariant in the original representation.
2. Choose a countable Hilbert basis, reindex it by a small countable type, and conjugate the action onto the resulting ℓ^2 carrier in the group's universe. The same-universe predicate then gives the polymorphic statement.

The pinned Mathlib already contains separability of spans (`TopologicalSpace.IsSeparable.span`), separability of closures (`TopologicalSpace.IsSeparable.closure`), completeness of closed spans (`Submodule.topologicalClosure.completeSpace`), and Hilbert-basis existence (`exists_hilbertBasis`).

At the pinned TauCeti revision, the continuous-representation library provides invariant restriction (`TauCeti.ContRepresentation.subrepresentation`), unitarity under restriction (`TauCeti.ContRepresentation.IsUnitary.subrepresentation`), and transport along an isometry (`TauCeti.ContRepresentation.IsUnitary.congr`).

Its compact-group lane also standardizes finite-dimensional carriers (`TauCeti.IrrepModel`) by an orthonormal basis, following the explicit finite-dimensional universe convention in its roadmap.

Neither pinned codebase assembles the countable invariant span, gives its Hilbert basis a small countable index, or transports the full property-(T) predicate across that model. Together these form the lowering seam.

They are more focused than generalizing the entire `Connes.spectral_criterion_unconditional` stack, but remain separate from closing the EJZK premise and from the present formalization.

12 Disclosure of AI use

12.1 Human direction and agent implementation

This is a human-directed, substantially agent-implemented formalization and exposition project. The division of work was concrete.

Human mathematical and proof direction

- selected Zhou's theorem and required completion of every Zhou-internal argument, with property (T) for $\mathrm{EL}_3(\mathbb{F}_2[t])$ as the single explicit external premise;
- fixed Zhou's Sections 2–7 as the proof order and required the overview to show one construction feeding factor equivalence, property (T), ICC, and nonisomorphism; and
- required comparison of the OpenAI, Zhou, and Anthropic proofs by the algebraic information each factor forgets, followed by the EJZK formalization outlook.

Human review and exposition contract

- required literature grounding, declaration-level provenance, public privacy boundaries, and Tau Ceti rubric review;
- required Linux Comparator verification with statement and axiom checks, confinement, Lean-kernel replay, and nanoda replay, while keeping technical acceptance distinct from mathematical audit;
- specified notes about the formalization rather than a restatement of the source papers, including proof architecture, mathematical weight, special designs, the threefold motivation, proof-family comparison, and unresolved boundaries; and
- reviewed rendered output and required corrections to headings, prose, diagrams, cross-references, navigation, privacy, and publication wording.

Agent implementation and review

- reconciled the literature and public Lean sources, adapted attributed proof blocks, implemented the Lean development, and refactored its public interfaces;

- ran literature, consumer-contract, provenance, semantic, and Tau Ceti reviews; built the Comparator, documentation, and rendering infrastructure; and repaired the resulting findings; and
- drafted and revised the mathematical notes, diagrams, comparisons, and formalization estimates under the human-specified structure and constraints.

The working cycle was: human proof and review contract ⟶ agent implementation ⟶ role-separated agent review and repair ⟶ human inspection and correction ⟶ approved landing.

Agent role separation is not independent human review; the project's audit ledger records the remaining review boundary. Public proof transfers are identified in the port map.

References

[Anthropic(2026)] Anthropic. 2026. Two Non-Isomorphic ICC Property (T) Groups with Isomorphic Group von Neumann Algebras. https://www-cdn.anthropic.com/files/4zrzovbb/website/1f7272990efcd274cbf92b6fdc14e3e280f59fdc.pdf (Cited in sections 1.2, 1.3.1, and 10.1)

[Ershov and Jaikin-Zapirain(2010)] Mikhail Ershov and Andrei Jaikin-Zapirain. 2010. Property (T) for Noncommutative Universal Lattices. *Inventiones Mathematicae* 179 (2010), 303–347. doi:10.1007/s00222-009-0218-0 (Cited in sections 1.1, 4, 5, 11.2, 11.3, and 11.4)

[Ershov et al.(2017)] Mikhail Ershov, Andrei Jaikin-Zapirain, and Martin Kassabov. 2017. Property (T) for Groups Graded by Root Systems. (2017). arXiv:1102.0031 https://arxiv.org/abs/1102.0031 (Cited in sections 4.5.1, 11.1, 11.2, 11.3, and 11.4)

[Kassabov(2005)] Martin Kassabov. 2005. Universal Lattices and Unbounded Rank Expanders. arXiv:math/0502237 https://arxiv.org/abs/math/0502237 (Cited in section 11.2)

[Nica(2019)] Bogdan Nica. 2019. On Bounded Elementary Generation for SLn over Polynomial Rings. arXiv:1901.00587 https://arxiv.org/abs/1901.00587 (Cited in sections 11.3 and 11.3)

[OpenAI(2026)] OpenAI. 2026. Ten Advances in Mathematics and Theoretical Computer Science. https://cdn.openai.com/pdf/ten-proofs-oai.pdf (Cited in sections 1.2, 1.3.1, 3.2, and 9.1)

[Zhou(2026)] Shuoxing Zhou. 2026. ICC Property-(T) Groups Without W*-Superrigidity. arXiv:2608.02327 [math.OA] https://arxiv.org/abs/2608.02327 (Cited in sections 1.1, 1.2, 1.3.1, 2.1.1, 3.1.1, 3.2, 5.1.1, 6.1.1, 7.1, 8.2, 8.3, 8.4, 9.1, and 11.3)

Alphabetical Index