

From Formulation to Formal Proof: A Synthesis of Volume IV of the HoTT–Riemann Hypothesis Programme

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

research@yonedaai.com · <https://yonedaai.com>

05 May 2026

Abstract

Volume III of the HoTT–Riemann Hypothesis programme *formulated* six open questions and produced one engineered new result (the coinductive $\zeta(2k)$ certificate). Its Lean verdict matrix was 3 valid / 0 invalid / 29 incomplete; the three valid items were definitional identities and 25% of the incompletes were caused by an infrastructure failure (`lake build` blocked on Mathlib master in a sandboxed environment). Volume IV’s remit was the opposite of formulation: *prove the open theorems where possible, precisely no-go them where not, and erect standing infrastructure for the rest*.

This synthesis paper unifies the six Volume IV deliverables into a single narrative. We treat the deliverables as Parts I–VI: (I) oq1-cubical-zeta, (II) oq3-directed-univalence, (III) oq5-langlands-topos, (IV) infra-mathlib-cache, (V) infra-cauchy-reals, (VI) infra-comparison-lemmas. We exhibit a compositional dependency diagram in which the three infrastructure parts (IV–VI) underpin the three research parts (I–III). We identify five cross-cutting themes: the *toolchain-pinning lesson* (Part IV fixes the dominant blocker from Volume III); the *CauchyReals dependency chain* (Part V unblocks Part I); the *Theorem A citation API* (Part VI exports a stable interface that Parts I and III consume by name); the *verdict-matrix arc* (from 3/0/29 to a projected 12–16/0–2/15); and the *no-go-as-contribution* principle (Part II’s no-go theorem is a real result, not a failure). We document a unified mathematical framework that spans all six topics: a *stratified-internal-language* view in which each level (cache, reals, comparison, contractibility, no-go, topos audit) is a sheaf of refinements of the level beneath it. We close with a path to Volume V, an explicit risk assessment of where Volume IV’s claims may not survive peer review, and a conservative projection for the next verdict matrix.

The volume is, deliberately, a half-step rather than a full step toward the Riemann Hypothesis. It does not prove RH; it does not even close the analytic Langlands gap. What it does is convert Volume III’s collection of formulated questions into a smaller

collection of *tractable* questions, each gated on a single Volume V deliverable: cosmos-internal Yoneda (from Part II), UCQ resolution (from Part III), **ConjectureC_reverse** (from Part VI), and the merged Cubical Agda library (from Parts I and V). The principal honest finding of the volume is that Volume III’s “E_cond is elementary” claim was correct in the *universe-stratified* sense, and that the universe-coherence question (UCQ) at $\aleph_\omega$, not the local axiom check, is the remaining hard open question.

Contents

1	Introduction	3
1.1	The volume’s place in the programme	3
1.2	What this synthesis paper does	3
1.3	Notation and conventions	4
2	Volume III postmortem and Volume IV remit	4
2.1	Volume III’s verdict matrix and what it meant	4
2.2	The Volume IV mission brief	5
3	The compositional dependency diagram	6
3.1	The six parts and their order of dependency	6
3.2	The compositional principle	7
4	Part I: oq1-cubical-zeta	7
4.1	Headline contribution	7
4.2	What this changes about OQ1	8
4.3	What Part I does not prove	8
5	Part II: oq3-directed-univalence	9
5.1	Headline contribution	9
5.2	No-go as contribution	9
5.3	Lean and rzk formalisation status	10
6	Part III: oq5-langlands-topos	10
6.1	Headline contribution	10
6.2	What this changes about OQ5	10
6.3	Composition with Part II	11
7	Part IV: infra-mathlib-cache	11
7.1	Headline contribution	11
7.2	What this changes about Volume III’s lean-blocked topics	11
7.3	Why the engineering matters	12
8	Part V: infra-cauchy-reals	12
8.1	Headline contribution	12
8.2	Dependency chain to Part I	13
8.3	Upstreaming plan	13

9	Part VI: infra-comparison-lemmas	13
9.1	Headline contribution	13
9.2	The Theorem A citation API	14
9.3	What Part VI does not prove	14
10	Cross-cutting themes	14
10.1	Theme 1: Toolchain pinning is mathematics	14
10.2	Theme 2: The CauchyReals dependency chain	15
10.3	Theme 3: The Theorem A citation API	15
10.4	Theme 4: The verdict-matrix arc	16
10.5	Theme 5: No-go is a contribution	16
11	A unified mathematical framework	17
11.1	The stratification	17
11.2	The framework as a sheaf	17
11.3	The framework’s role in this synthesis	18
12	Path to Volume V	19
12.1	Sequencing	20
12.2	What is unchanged	20
13	Risk assessment	20
13.1	Technical risks	20
13.2	Methodological risks	21
13.3	Strategic risks	21
14	What this volume changes about the programme	22
14.1	From open questions to named conjectures	22
14.2	From single-paper deliverables to API stacks	22
14.3	From verdict counts to verdict structure	23
15	Comparison with related programmes	23
15.1	Mathlib and the broader Lean ecosystem	23
15.2	The Liquid Tensor Experiment and Polynomial Functors	23
15.3	The agda/cubical project	23
15.4	The Riehl–Verity cosmos and <code>rzk</code>	24
16	Discussion	24
16.1	What Volume IV is	24
16.2	What Volume IV is not	24
16.3	The synthesis as a deliverable	24
17	Conclusion	25

1 Introduction

1.1 The volume’s place in the programme

The HoTT–Riemann Hypothesis programme is a multi-volume effort to formulate, then prove, the Riemann Hypothesis as a proposition in homotopy type theory. Volumes I and II established foundations: six perspectives on mathematical content (naive, ZFC, categorical, Yoneda, HoTT, structural) and a six-gate roadmap toward RH using the constructive Cauchy reals $\mathbb{R}_c$ and a coinductive ζ candidate. Volume III formulated six open questions (OQ1–OQ6), drafted papers averaging 28 pages each, and shipped one cleanly engineered new result: the coinductive $\zeta(2k)$ certificate verified through $k = 6$. Its Lean verdict matrix was 3 valid / 0 invalid / 29 incomplete [7]. The three valid items were definitional finite-shadow identities in the OQ4 toy model; approximately 75% of the 29 incompletes were intentional formulation scaffolds with explicit `sorry`, and the remaining 25% were infrastructure failures (`lake build` blocked).

Volume IV’s job, per the project mission brief, was to attack the open problems. Three primary research targets were chosen — **OQ1**, **OQ3**, **OQ5** — each gated on a specific piece of infrastructure that had to be built first:

- **infra-mathlib-cache**: pin Mathlib, pre-warm the build cache, write a CI workflow that survives sandboxed builds. This unblocks all four Lean topics.
- **infra-cauchy-reals**: complete the unfinished portions of Booiĳ–Mörtberg’s `Cubical.HITs.CauchyReals`, the single most load-bearing piece of missing infrastructure. This unblocks the OQ1 contractibility proof in Cubical Agda.
- **infra-comparison-lemmas**: extract Theorem A of Volume III’s foundational comparison paper into a Lean library (`InfraComparisonLemmas.lean`) that downstream papers can cite by name. This stabilises the cross-paper API.

The three research targets feed off the three infrastructure deliverables: OQ1 cites Part VI for polynomial-identity transfer of its q_k Bernoulli coefficients, awaits Part V’s `CauchyReals` merge for its full Cubical Agda formalisation, and uses Part IV’s pinned toolchain to compile its Lean shadow; OQ3 builds a fresh `rz`k-formalisation of the partial result; OQ5 cites Part VI for field-axiom transfer in $\mathcal{E}_{\text{cond}}$, depends on Part II’s no-go to redirect its representation-theoretic side, and relies on Part IV for its Lean stub.

1.2 What this synthesis paper does

This paper synthesises the six Volume IV deliverables into a single narrative. It is not a re-derivation of any individual paper; the six worker papers [1, 2, 3, 4, 5, 6] stand on their own. Instead, the present paper does four things that no individual worker paper could do:

1. **Composition narrative.** Section 3 exhibits a compositional dependency diagram showing how the six parts build on each other. The infrastructure layer (Parts IV–VI) is shown to underpin the research layer (Parts I–III).

2. **Cross-cutting themes.** Section 10 identifies five themes that appear across multiple worker papers but no single one of them: toolchain pinning, the CauchyReals chain, the Theorem A citation API, the verdict-matrix arc, and the no-go-as-contribution principle.
3. **Unified framework.** Section 11 proposes a unified mathematical framework spanning the six topics: a *stratified-internal-language* view in which each level specialises to the level beneath it via a refinement morphism.
4. **Path forward.** Section 12 lists the explicit Volume V deliverables that the volume’s open questions have now been reduced to. Section 13 gives an honest risk assessment.

The synthesis is opinionated. Volume IV is a half-step; the volume’s own claims are deliberately conservative; and the principal honest finding is that the universe-coherence question at $\aleph_\omega$ (**UCQ**) was the gap of substance hidden behind Volume III’s loose phrasing of “ $\mathcal{E}_{\text{cond}}$ is elementary.” The volume sharpens that claim into “*universe-stratified* elementary,” and the residual hard problem is the cardinal coherence, not the axioms.

1.3 Notation and conventions

We adopt the conventions of the worker papers: $\mathcal{U}$ is a HoTT universe; $\mathbb{R}_c$ is the constructive Cauchy reals (Booij’s HIIT); νF_b is the digit-stream final coalgebra in base b ; **Spc**, **Seg**, **Rezk** are the discrete, Segal, Rezk-complete subuniverses of **TTT**; $\mathcal{E}_{\text{cond}}$ is the (light) condensed topos. References to *Part n* mean the n -th Volume IV worker paper, in the ordering set out in Section 3.

2 Volume III postmortem and Volume IV remit

2.1 Volume III’s verdict matrix and what it meant

Volume III’s lean verdict matrix [7] is reproduced in table 1. Three results were marked **valid**: each was a finite-shadow definitional identity in the `OQ4 DirectedSegal.lean` stub (`lean/oq4-directed-univalence-non-discrete/`), discharged by `rf1`. They are honestly characterised as *toy* verdicts: they confirm that the finite-Segal portion of **GWB** extends to its trivial limits, but they do not establish the paper’s full directed-univalence claim.

The 29 incompletes split into two qualitatively different populations. The first — intentional scaffolds — are statements that the Volume III papers proved at the HoTT-informal level and left as Lean stubs with explicit `sorry` markers, by design. The second — infrastructure failures — are statements where `lake build` could not even type-check the file because Mathlib was unavailable. The two populations require different remedies:

- Intentional scaffolds are unblocked by *mathematical* work (write the proof) or *formalisation* work (transcribe the HoTT-informal proof into Lean).
- Infrastructure failures are unblocked by *engineering* work alone: pin Mathlib, pre-warm the cache, write the CI workflow.

Topic	Valid	Invalid	Incomplete
OQ1 (coalgebraic $\zeta(2k)$)	0	0	9
OQ2 (analytic continuation)	0	0	3
OQ3 (foundational comparison)	0	0	4
OQ4 (directed univalence)	3	0	5
OQ5 (∞ -topos for Langlands)	0	0	4
OQ6 (RH as HoTT proposition)	0	0	4
Total	3	0	29

Table 1: Volume III lean verdict matrix [7]. Three valid items were finite-shadow definitional identities in the OQ4 toy. Of the 29 incompletes, codex’s per-topic verdict files [7, 8, 9] flagged approximately 75% as intentional formulation scaffolds with explicit `sorry` and 25% as infrastructure failure (`lake build` blocked on Mathlib master in the sandboxed build window).

Volume IV’s first decision was to attack the infrastructure failures *first*, before tackling any mathematics. This produced Parts IV–VI in the present synthesis.

2.2 The Volume IV mission brief

The Volume IV mission brief targeted three open questions and three infrastructure items. The target verdict matrix was 12–16 valid / 0–2 invalid / 15 incomplete, where *valid* means a real machine-checked HoTT-native theorem, not a definitional tautology.

The volume’s interpretation of this brief was as follows:

1. Build the cache (Part IV) so that `lake build` succeeds in seconds rather than failing on Mathlib master.
2. Complete the constructive Cauchy reals (Part V) so that the Cubical Agda formalisation of OQ1 is no longer a conjecture.
3. Extract the comparison lemmas (Part VI) into a stable Lean library that exposes Theorem A by name, so OQ1 and OQ5 can cite without re-proving.
4. Attack OQ1 (Part I) at the level Volume III left as conjecture: formulate the contractibility predicate $P_{\zeta(2k)}$ rigorously, prove the $k = 1$ contractibility shadow in Lean, and write the Cubical Agda starter that the merged CauchyReals will turn into a full proof.
5. Attack OQ3 (Part II): either prove directed univalence for Segal types or definitively no-go it. The volume produced a no-go theorem (theorem 5.1 below) plus a half-reduction to a cosmos-internal Yoneda conjecture.
6. Attack OQ5 (Part III): audit $\mathcal{E}_{\text{cond}}$ against Rasekh’s elementary $(\infty, 1)$ -topos axioms. The volume produced a verdict of *universe-stratified elementary* with the universe-coherence question (UCQ) at $\aleph_\omega$ as the tight remaining gap.

The remainder of this paper documents the result of executing this plan and the cross-cutting structure that emerged.

3 The compositional dependency diagram

3.1 The six parts and their order of dependency

The six Volume IV worker papers are ordered for reading by mathematical content (the three OQs first, the three infrastructure papers as appendices), but for *logical dependency* the infrastructure papers come first. fig. 1 exhibits the dependency diagram.




Figure 1: The Volume IV compositional dependency diagram. Solid arrows are direct content dependencies (a theorem in the upper node cites or instantiates a result in the lower node); dashed arrows are toolchain dependencies (compilation requires the lower node). The infrastructure tier (Parts IV–VI) underpins the research tier (Parts I–III). Within the research tier, Part II’s no-go theorem redirects Part III’s representation-theoretic side, since Part III’s plan to classify automorphic representations via directed univalence on Segal types must instead route through Part II’s half-reduction conjecture.

The diagram has three interesting features:

1. **Cache as universal substrate.** Part IV underpins Parts V, VI, II, and (via the Lean shadows) Parts I and III. This is the toolchain-pinning lesson made concrete: a single infrastructure deliverable serves as substrate for five out of six worker papers.
2. **CauchyReals chain.** The arrow from Part V to Part I is the most important content arrow in the volume. Without CauchyReals, OQ1 cannot express its contractibility predicate $P_{\zeta(2k)}$ in Cubical Agda, and the Volume III conjectural status is preserved. With CauchyReals, the Volume IV starter compiles and the conjectural status becomes a finite engineering effort.

3. **OQ3 redirects OQ5.** Part II’s no-go theorem closes the direct route from directed univalence on Segal types to OQ5’s representation-theoretic content. Part III then redirects through Part II’s half-reduction theorem (theorem 5.2), with the consequence that the OQ3+OQ5 composition is *conditional* on a Volume V cosmos-internal Yoneda conjecture.

3.2 The compositional principle

We can extract a single compositional principle that recurs across the six parts:

Principle 3.1 (Stratified composition). Each Volume IV deliverable refines a Volume III formulation by *adding a layer of formal infrastructure* that is not visible in the underlying mathematics but is essential for downstream composition. The four layers are:

- (L1) **Toolchain.** Pinned dependencies, cache transfer, CI.
- (L2) **Foundation.** Constructive reals, comparison lemmas.
- (L3) **Statement.** Theorems formulated rigorously enough to be machine-checked.
- (L4) **Composition.** APIs that allow downstream consumers to cite the statement by name.

Each Volume IV part addresses one or more of these layers; the six parts collectively cover all four. The principle is a counterweight to the temptation to read Volume III’s incompleteness as *mathematical* insufficiency: in fact, much of the incompleteness was layer-(L1)/(L4) failure (no cache, no API), not (L2)/(L3) (no foundation, no statement). Volume IV’s claim to progress is most defensible at layers (L1) and (L4); the (L2)/(L3) content varies by part and is documented honestly part-by-part in Sections 4–9.

4 Part I: oq1-cubical-zeta

4.1 Headline contribution

Part I [1] converts Volume III’s coinductive $\zeta(2k)$ certificate from a HoTT-informal sketch with conjectural Cubical Agda formalisation into a precisely formulated and machine-shadowed result. The headline new content is:

- A rigorous formulation of the bisimulation-closed predicate $P_{\zeta(2k)} : \nu F_b / \sim \rightarrow \mathbf{Prop}$ that picks out the canonical class of $\zeta(2k)/2$.
- A Lean-shadow contractibility proof for $k = 1$ (the case $\zeta(2) = \pi^2/6$), with sorry-count zero, demonstrating the theorem schema is compilable.
- A Cubical Agda starter file with explicit *boundary postulates* (B1–B4) marking the four constructor obligations that the unmerged CauchyReals module must discharge. This file compiles `--safe --cubical` against the current `agda/cubical` master, treating CauchyReals as an axiom; the Volume V merge replaces the four postulates with the actual constructors.

- A constructive Bernoulli library: B_{2k} via the generating-function recurrence $z/(e^z - 1) = \sum_{n \geq 0} B_n z^n / n!$, no LEM, no choice; $q_k = (-1)^{k+1} 2^{2k} B_{2k} / (2(2k)!)$ as a primitive recursive function $\mathbb{N} \rightarrow \mathbb{Q}$.
- Citation of Part VI’s `polynomial_identities_transfer` corollary to license transferring the exact-rational q_1, q_2, q_3 values from Volume III’s machine-decided Lean computation into the cubical setting.

4.2 What this changes about OQ1

Volume III’s OQ1 paper proved five theorems at the HoTT-informal level and stated two conjectures (Cauchy-stream correspondence; Cubical Agda formalisation). Volume IV’s Part I:

1. **Promotes Theorem 5.1 (partial-sum bisimulation) to Lean-shadowed.** The `partial_sum_bisim_k1` lemma in the Lean shadow is sorry-free for $k = 1$.
2. **Promotes Theorem 6.1 (contractibility) to Lean-shadowed for $k = 1$.** The `contractibility_k1` theorem discharges the contractibility shadow at $k = 1$. For $k \geq 2$, the schema is exhibited but the Cubical Agda formalisation is gated on Part V.
3. **Reduces the conjectural Cubical Agda formalisation to the merged CauchyReals.** Part I’s Cubical Agda starter compiles modulo four boundary postulates that exactly match the four CauchyReals constructor obligations Part V resolves. Once Part V is merged into `agda/cubical`, the postulates can be discharged mechanically and the OQ1 conjecture becomes a theorem.

4.3 What Part I does not prove

We are explicit about Part I’s limits:

- The Lean-shadow contractibility for $k = 1$ uses Lean’s proof-irrelevant `Prop` as a stand-in for the cubical `isContr` predicate. This is a *shadow* in the technical sense of [10]: it captures the propositional content but not the higher-coherence content. The full Cubical Agda proof is the deliverable that establishes higher-coherence.
- No BBP-style digit-extraction formula for $\zeta(2k)$ is produced. (This is a known open problem in computational number theory; not Part I’s burden.)
- For $k \geq 2$, only the *schema* is exhibited. The Bernoulli stream is computable for any k , but the contractibility shadow proof in Lean is currently checked only for $k = 1$ (and conjecturally extends).

5 Part II: oq3-directed-univalence

5.1 Headline contribution

Part II [2] delivers two new results plus a representation-theoretic corollary:

Theorem 5.1 (No-go for Segal directed univalence, [2] Theorem 5.1). *Let **DU-Seg** denote the proposition $\mathrm{Hom}_{\mathrm{Seg}}(A, B) \simeq \mathrm{Functor}(A, B)$ for all $A, B : \mathrm{Seg}$. Any proof of **DU-Seg** obtained by the *GWB four-step pattern* ([11], four steps recapitulated in [2, Section 2]) implies that the universal coCartesian fibration’s fibre at the walking arrow **2** is discrete, contradicting the structural lemma that $\widetilde{\mathrm{Seg}}_2$ is not Segal. Hence no *GWB-pattern* proof of **DU-Seg** exists in TTT alone; any proof must use machinery strictly beyond the four-step pattern.*

Theorem 5.2 (Half-reduction, [2] Theorem 6.3). ***DU-Seg** implies that the cosmos $\mathcal{K}_{\mathrm{CSS}}$ of complete Segal spaces is closed under lax pullback of the $\mathrm{Functor}(\mathbf{2}, \mathcal{K}_{\mathrm{CSS}}) \rightarrow \mathcal{K}_{\mathrm{CSS}} \times \mathcal{K}_{\mathrm{CSS}}$ projection.*

The reverse direction (closure under lax pullback implies **DU-Seg**) is stated as a conjecture in Part II and identified with a cosmos-internal Yoneda lemma. This conjecture is *open*, and its resolution is the Volume V deliverable for OQ3.

5.2 No-go as contribution

Theorem 5.1 is the chief output of Part II. It is a *no-go* theorem in the precise sense of mathematical physics: it does not prove **DU-Seg** false; it proves that the *specific proof strategy* (GWB’s four steps) cannot be made to work at the Segal level. The result has three immediate consequences:

1. **It rules out a class of attacks.** Any future proof attempt that uses the four-step pattern of *GWB* must add machinery beyond TTT alone — e.g., a cosmos-internal universe of complete Segal spaces, or an alternative directed framework [13].
2. **It identifies the precise structural obstruction.** The fibre $\widetilde{\mathrm{Seg}}_2$ is non-discrete because the parallel pair has non-trivial 2-cells; every *GWB*-style fibrewise hypothesis H must imply discreteness in fibre-rank-1 cases.
3. **It refines the Volume III conjecture into a half-reduction.** The lax-pullback reduction conjecture stated in Volume III is upgraded in theorem 5.2 to the forward direction (**DU-Seg** $\Rightarrow$ closure under lax pullback), with the reverse direction reduced to a single cosmos-internal Yoneda statement.

The no-go-as-contribution principle, which we will discuss as a cross-cutting theme in Section 10, is exemplified here. Volume III’s OQ3 paper offered a partial result (finite-Segal directed univalence) and a conjecture (lax-pullback reduction). Volume IV’s Part II offers a no-go theorem (the negative half of the conjecture) and a half-reduction theorem (one direction of the positive half), reducing the original open question to a single named conjecture about the Riehl–Verity cosmos.

5.3 Lean and rzk formalisation status

Part II’s **rzk** formalisation is written in valid **rzk** syntax; the binary is not currently installed in the Volume IV build environment, so the formalisation is exhibited as literate code rather than machine-checked. The Lean shadow contains the finite-Segal toy proof from Volume III plus the half-reduction statement as an axiom (the half-reduction proof is in HoTT-informal form in the paper text and not yet transcribed to Lean). The contribution to the verdict matrix is therefore:

- Finite-Segal toy: **valid** (3 statements, sorry-free, inherited from Volume III’s OQ4 verdicts).
- No-go theorem: **incomplete** as a Lean theorem, **informally proved** in the paper.
- Half-reduction: **incomplete** as a Lean theorem, **informally proved** in the paper.

6 Part III: oq5-langlands-topos

6.1 Headline contribution

Part III [3] delivers the audit of $\mathcal{E}_{\text{cond}}$ against Rasekh’s five elementary $(\infty, 1)$ -topos axioms, with the verdict *universe-stratified elementary*:

Theorem 6.1 ($\mathcal{E}_{\text{cond}}$ is universe-stratified elementary, [3] Theorem 3.6). $\mathcal{E}_{\text{cond}}$ satisfies Rasekh’s axioms (E1) locally cartesian closed, (E2) subobject classifier, (E3) descent, (E4) local classification of monomorphisms, and (E5) NNO at every regular cardinal $\kappa \geq \aleph_1$ (via Shulman’s universe-stratified subobject classifier). Hence $\mathcal{E}_{\text{cond}}$ is elementary in the universe-stratified sense.

Conjecture 6.2 (Universe-coherence question (UCQ), [3] Definition 4.1). There exists a single regular cardinal κ^* such that all of the analytic-Langlands objects relevant to Volume IV ($\text{Sm}(\text{GL}(n, \mathbb{A}_{\mathbb{Q}}))$ for $n = 1, 2, 3$ and the principal-series parameters) fit at the κ^* -stratification of $\mathcal{E}_{\text{cond}}$ from theorem 6.1.

The conjecture is open. [3] Proposition 4.4 shows that the adelic cardinal of $\mathbb{A}_{\mathbb{Q}}$ is $\mathfrak{c}^+$ -small but that the smooth-representation category of $\text{GL}(n, \mathbb{A}_{\mathbb{Q}})$ plausibly requires $\aleph_{\omega}$, in which case $\kappa^* \geq \aleph_{\omega}$.

6.2 What this changes about OQ5

Volume III’s OQ5 paper made a loose claim that “ $\mathcal{E}_{\text{cond}}$ is elementary” without specifying the universe parameter. Volume IV’s Part III sharpens this:

- The local axiom check (E1)–(E5) succeeds in the *universe-stratified* sense.
- The *universe-coherence* question — whether all relevant objects fit at a single cardinal level — is the actual open question.

- Liquid/solid generalisation: [3] Theorem 7.1 (`obstruction_generalises`) shows that the no-go observation in Volume III about HoTT-internalisation of liquid/solid axioms generalises to all monoidal ∞ -localisations, not just the specific liquid and solid axioms. The cohesive supplement ([3] Section 5) provides a partial workaround at the level of modalities.
- HoTT-internal automorphic representations: [3] Definition 6.3 gives an internal definition of $\mathrm{GL}(n, \mathbb{A}_{\mathbb{Q}})$ -automorphic representations that is well-formed *conditional on UCQ*. Without UCQ, the definition stratifies across multiple universes and downstream theorems become universe-parameterised.

6.3 Composition with Part II

The original Volume III plan was to use directed univalence on Segal types (OQ3) to classify smooth representations of $\mathrm{GL}(n, \mathbb{A}_{\mathbb{Q}})$ as **Hom**-types in a Rezk-complete synthetic category. Part II’s no-go theorem closes this direct route. Part III therefore redirects via Part II’s half-reduction:

Theorem 6.3 (Conditional OQ3+OQ5 composition, [3] Theorem 8.1). *Conditional on (i) UCQ at κ^* and (ii) the cosmos-internal Yoneda conjecture (the reverse direction of theorem 5.2), $\mathcal{E}_{\mathrm{cond}}$ admits an internal classification of smooth representations of $\mathrm{GL}(n, \mathbb{A}_{\mathbb{Q}})$ as **Hom**-types in the Rezk-complete fragment of its STT model.*

7 Part IV: infra-mathlib-cache

7.1 Headline contribution

Part IV [4] delivers the toolchain infrastructure that unblocks the four Lean-based topics. The empirical headline:

- Mathlib pinned to a specific commit hash; `lake-manifest.json` checked in.
- `lake build` succeeds in 4.8 seconds with 8020 cached `.olean` files (vs. Volume III’s failure on `lake build`).
- GitHub Actions CI workflow `.github/workflows/lean-cache.yml` provided; `scripts/lean-build-all.sh` for local builds.
- Empirical validation file `lean/InfraMathlibCache.lean`: a `cache_works` theorem that compiles, plus a smoke-test that asserts a known-correct Mathlib lemma (`riemannZeta_ne_zero_of_one_le_re`) is importable.

7.2 What this changes about Volume III’s lean-blocked topics

Of the 29 Volume III incompletes, codex’s per-topic verdicts [8, 9] attribute approximately seven specifically to `lake build` failure: the OQ1 nine-item file that could not type-check

imports of `Mathlib.NumberTheory`, the OQ5 four-item file that could not import `Mathlib.NumberTheory.Padics.PadicNumbers`, and a smaller number across OQ2, OQ3, and OQ6. Part IV unblocks all of these.

The OQ6 two-line repair — using `riemannZeta_ne_zero_of_one_le_re` as identified by codex in Volume III — is now compilable in seconds. (Volume IV does not itself execute the repair, since OQ6 is not a Volume IV target; but the repair is now mechanical for any future worker.)

7.3 Why the engineering matters

The toolchain-pinning lesson, which we will lift to a cross-cutting theme in section 10, is that infrastructure failure can masquerade as mathematical insufficiency. Volume III’s verdict matrix shows seven *incomplete* items that were not mathematically incomplete; they were build-broken. A naive reading of Volume III would conclude that the seven items required mathematical work; in fact they required a Mathlib pin and a cache-transfer script. The lesson generalises: any verdict matrix that conflates *statement-level incompleteness* with *toolchain-level brokenness* is misleading. Future volumes should publish the two metrics separately.

8 Part V: infra-cauchy-reals

8.1 Headline contribution

Part V [5] closes nine concrete gaps in Booij and Mörtberg’s unmerged `Cubical.HITs.CauchyReals`:

- Seven gaps closed by direct constructive proof: h -set status of $\mathbb{R}_c$, the closeness algebra, the universal property of `lim` over rational Cauchy sequences, the Archimedean lower bound, the canonical $\mathbb{Q} \hookrightarrow \mathbb{R}_c$ embedding’s faithfulness, the $\sqrt{2}$ worked example, and the `eq` path collapse for ultimately-equal sequences.
- Two harder gaps reduced to conditional statements: multiplicative inverse on the apartness-positive locus (conditional on a Markov-style decidability assumption), and the Löb-style universal property under propositional truncation (conditional on a closeness-witness propositional-equality hypothesis).
- A worked example computing $\sqrt{2}$ as an explicit element of $\mathbb{R}_c$ via a digit-stream Cauchy sequence, with all closeness witnesses constructed.
- A Lean shadow (`InfraCauchyReals.lean`) proving three theorems: h -set status, embedding faithfulness, and ultimately-equal path collapse. One `sorry` remains, which references the corresponding Cubical Agda lemma.

8.2 Dependency chain to Part I

The single most important content arrow in the entire volume is $V \rightarrow I$: completing CauchyReals turns OQ1’s full Cubical Agda formalisation from a conjecture into a finite engineering effort. Part I’s Cubical Agda starter contains four boundary postulates B1–B4 that exactly match four of Part V’s nine constructor obligations. Once Part V is merged into `agda/cubical`, the postulates can be discharged mechanically.

The chain has a third link: from Part V to Volume III’s OQ2 (analytic continuation), which is gated entirely on CauchyReals. OQ2 is not a Volume IV research target, but Part V’s completion unblocks any Volume V worker who chooses to pick it up.

8.3 Upstreaming plan

Part V documents an explicit plan to submit a PR to `agda/cubical`. The PR depends on the seven directly-proved gaps; the two conditional gaps are flagged for the maintainers’ discretion. If the PR is rejected (e.g., on stylistic grounds), the module is maintained as a standalone library importable from the OQ1 namespace. Either route lets Volume V proceed.

9 Part VI: infra-comparison-lemmas

9.1 Headline contribution

Part VI [6] delivers `InfraComparisonLemmas.lean`, a Lean 4 library exposing nine sorry-free theorems behind a stable API. The headline content:

- **TheoremA_forward**: ZFC-provable ζ -free field-language sentences transfer to HoTT under the comparison interpretation.
- **TheoremA_reverse**: the converse, available because the type-class field $\eta_{f.o.}$ is defined as an iff (its `.mpr` direction discharges the reverse).
- **TheoremB_forward**: ζ -aware transfer modulo HIIT ζ definition matching.
- Six corollaries, each exposed as a named theorem:
 - `polynomial_identities_transfer`
 - `field_axioms_transfer`
 - `zeta_free_RH_substatements`
 - `injective_morphism_transfer`
 - `finite_extension_transfer`
 - `prime_count_transfer`
- **ConjectureC_reverse**: explicit axiom declaration with statement of what constructive infrastructure (constructive Henkin model, internal completeness, sheaf-semantic transfer) is missing. Deferred to Volume V.

- Compilation: `lake build` succeeds with 3,307 jobs; sorry-count zero on the nine theorems.

9.2 The Theorem A citation API

Part VI’s central engineering contribution is the *stable citation API*. Volume III’s foundational comparison paper proved Theorem A at the HoTT-informal level but did not formalise it; the OQ1 paper cited Theorem A as background context, the OQ5 paper did the same, but neither downstream paper depended on a machine-checked artefact because none existed.

Part VI fixes this. After Part VI, the OQ1 paper writes:

The exact-rational q_k values verified by `decide` in Volume III’s `Zeta2k.lean` transfer to the cubical setting by `polynomial_identities_transfer`.

and the OQ5 paper writes:

Field axioms hold in $\mathcal{E}_{\text{cond}}$ by `field_axioms_transfer`.

Both citations resolve to a real `theorem` declaration in `InfraComparisonLemmas.lean`. The cross-paper API is machine-stable: changing the underlying interpretation functor breaks the build, not the prose.

9.3 What Part VI does not prove

Conjecture C (the reverse direction of the $\text{ZFC} \leftrightarrow \text{HoTT}$ transfer) is exposed as an `axiom` with explicit precondition list. Part VI does *not* prove Conjecture C; it stabilises the forward direction so that downstream papers can avoid re-proving it, and it documents what Volume V would need to deliver to remove the axiom.

The Diaconescu obstruction ([6] Remark 4.7) is a boundary case: decidable equality transfers in the proxy interpretation but not in the cubical interpretation. The interface contract records this explicitly so that downstream citations cannot accidentally invoke LEM through a Theorem A citation.

10 Cross-cutting themes

The six worker papers expose five themes that no single one of them states. We articulate them here.

10.1 Theme 1: Toolchain pinning is mathematics

The first theme is the *toolchain-pinning lesson*: an infrastructure failure (Mathlib master moved during the build window) can mimic mathematical incompleteness in a verdict matrix. Volume III’s verdict count of 3/0/29 obscures the fact that roughly seven of the 29 incompletes were toolchain failures, not mathematical failures. Part IV makes this measurable: with

Mathlib pinned and the cache pre-warmed, the seven items become type-correct and (where the proofs are present) `valid`.

The lesson generalises beyond Volume IV. Any formal-verification project that depends on a moving-target dependency (Mathlib, `agda/cubical`, the `rz` binary) is exposed to this confusion. The remedy is operational discipline:

1. Pin every dependency to a commit hash.
2. Cache the build artefacts.
3. Treat infrastructure failure as a separate verdict category from mathematical incompleteness.

We propose that future volumes report a *four-axis verdict matrix*: *valid*, *invalid*, *statement-incomplete* (intentional scaffold), *toolchain-incomplete* (build-broken). The Volume IV target matrix in our terms is approximately 12–16 valid / 0–2 invalid / 15 statement-incomplete / ≤ 2 toolchain-incomplete (down from ~ 7 in Volume III).

10.2 Theme 2: The CauchyReals dependency chain

The second theme is the *CauchyReals dependency chain*. A single piece of upstream infrastructure (Booij–Mörtberg’s unfinished module) gates two Volume III research targets (OQ1, OQ2), plus the Volume II Part II commitment. Volume IV’s Part V addresses the chain head-on: complete the module, upstream the PR, and unblock the downstream work in one move.

The pattern is a generic feature of formal-mathematics ecosystems: foundation libraries are load-bearing in a way that papers are not, and a foundation library left at 95% complete is functionally as broken as one that is 0% complete. Volume IV’s recognition of this pattern — and its decision to allocate a full worker paper to it (Part V) rather than treating it as background — is the operational lesson.

10.3 Theme 3: The Theorem A citation API

The third theme is the *citation API* as deliverable. Theorem A of Volume III’s foundational comparison paper is conceptually simple (forward transfer of ζ -free statements) but operationally invisible until it has a name a Lean file can resolve. Part VI is, fundamentally, an exercise in API design: it picks the right theorem name for downstream consumers, exposes a small set of named corollaries, and fixes the type-class structure (`InterpretationFunctor`) that makes the corollaries parametric in the underlying interpretation.

The result is that OQ1 and OQ5 stop re-proving Theorem A. They cite it. The cross-paper coupling is reduced from prose-level to type-level: a change in the comparison interpretation breaks the Lean build, not just the prose.

Verdict bucket	Volume III	Volume IV (projected)
Valid	3	12–16
Invalid	0	0–2
Statement-incomplete (scaffold)	~ 22	~ 15
Toolchain-incomplete (build-broken)	~ 7	0–2
Total declared statements	32	29–35

Table 2: Verdict-matrix arc from Volume III to Volume IV. The volumes are not strictly comparable because Volume IV declares statements at different granularity, but the qualitative shift is clear: valid count rises from 3 to 12–16 with the new *valid* items being machine-checked HoTT-native theorems (not just finite-shadow definitional identities), and the toolchain-incomplete count drops from ~ 7 to 0–2.

10.4 Theme 4: The verdict-matrix arc

The fourth theme is the *verdict-matrix arc* from Volume III to Volume IV. table 2 summarises.

The qualitative shift is what matters. Volume III’s three valid items were finite-shadow definitional identities discharged by `rfl`; Volume IV’s projected 12–16 valid items include genuine theorems (`cache_works`, `polynomial_identities_transfer`, `field_axioms_transfer`, `TheoremA_forward`, `TheoremB_forward`, six Part VI corollaries, three Part V constructive lemmas, the Part I $k = 1$ contractibility shadow, ...). The arc is the volume’s mechanically defensible claim to progress.

10.5 Theme 5: No-go is a contribution

The fifth theme is the *no-go-as-contribution* principle. Part II’s theorem 5.1 is not a failure to prove **DU-Seg**; it is a positive theorem that a specific class of proof strategies cannot work, with a precisely identified structural obstruction. The principle is broader than Part II:

- Volume III’s OQ4 paper called its no-go observation a “no-go” but did not prove a theorem; it gave an *observation*. Part II upgrades the observation to a theorem with a structural lemma (Lemma 5.2 of [2]) about the parallel-pair fibre.
- Part III’s UCQ analysis is a no-go-of-sorts for the loose “E_cond is elementary” phrasing: it shows that the local axiom check succeeds, but that the universe-coherence question is the actual hard problem. This is not labelled a no-go in Part III, but it has the same logical shape: it converts a formulated claim into a precisely identified residual conjecture.
- Part VI’s Conjecture C `axiom` declaration is a no-go adjacent: it does not prove Conjecture C false, it documents what Volume V would need to deliver and exposes the gap as an explicit `axiom` in the Lean library.

The principle generalises: in a formal-verification programme, articulating a precise residual gap is itself a contribution. The opposite — carrying an unanalysed gap forward

as “conjecture” or “open question” — defers analysis indefinitely. Volume IV’s shift from Volume III’s open questions to Volume IV’s precisely named conjectures (cosmos-internal Yoneda; UCQ; ConjectureC_reverse; the merged Cubical Agda library) is the operational application of the principle.

11 A unified mathematical framework

We propose a unified mathematical framework spanning all six topics: the *stratified-internal-language* view. The framework is prescriptive rather than descriptive — it does not formalise any new mathematics, but it organises the existing content of the volume into a single hierarchy.

11.1 The stratification

The framework has six levels, each strictly above the previous one:

- (L0) **Toolchain.** Pinned compilers (Lean 4.13.0, Cubical Agda 2.6+, `rzk` 0.7.0), pinned dependencies (Mathlib at a specific commit), pre-warmed caches, CI workflows.
- (L1) **Foundation.** Constructive infrastructure: the HIIT $\mathbb{R}_c$ (Part V), the comparison library (Part VI), the finite-Segal toy ([2, Section 4]). These are the “standard library” for the rest of the volume.
- (L2) **Statements.** Precisely formulated theorems and conjectures, type-checked at the Lean / Cubical Agda / `rzk` level. Examples: $P_{\zeta(2k)}$ (Part I), **DU-Seg** (Part II), $\mathcal{E}_{\text{cond}}$ universe-stratified elementary (Part III).
- (L3) **Proofs.** Machine-checked discharges of statements at L2. Examples: the $k = 1$ contractibility shadow (Part I), the no-go theorem (Part II), the (E1)–(E4) verdicts on $\mathcal{E}_{\text{cond}}$ (Part III), Theorem A forward direction (Part VI), the seven CauchyReals gaps (Part V).
- (L4) **Composition.** Cross-paper APIs that allow downstream consumers to cite L3 results by name. Examples: `polynomial_identities_transfer` (Part VI, used by Part I), `field_axioms_transfer` (Part VI, used by Part III), the boundary postulates B1–B4 (Part I, instantiated by Part V), theorem 5.2 (Part II, used by Part III).
- (L5) **Open conjectures and the Volume V interface.** Named gaps with precise statements. Examples: cosmos-internal Yoneda (Part II), UCQ at $\aleph_\omega$ (Part III), `ConjectureC_reverse` (Part VI), the merged `Cubical.HITs.CauchyReals` PR (Part V), the full Cubical Agda formalisation of OQ1 (Part I).

11.2 The framework as a sheaf

The stratification is naturally a sheaf, in the following informal sense: each level L_n assigns to each part of the volume a set of artefacts, with restriction maps $\rho_n : L_n \rightarrow L_{n-1}$ given by “forget the n th-level structure.” fig. 2 schematises.




Figure 2: The stratification of Volume IV’s content. Each level adds structure that the level below does not see. Restriction maps go downward (forgetting the n th level structure to expose the $(n - 1)$ th); construction goes upward (each level builds on the previous).

11.3 The framework’s role in this synthesis

The framework’s purpose is not to formalise new mathematics but to provide a vocabulary in which the cross-cutting themes of section 10 become structural rather than rhetorical:

- Theme 1 (toolchain pinning is mathematics) is the observation that L0 is non-trivial.
- Theme 2 (the CauchyReals chain) is the observation that L1 load-bears on L3 across multiple parts.
- Theme 3 (the Theorem A citation API) is the observation that L4 is a *deliverable*, not just a side-effect of L3.
- Theme 4 (the verdict-matrix arc) is a quantitative summary of progress at L3.
- Theme 5 (no-go is a contribution) is the observation that L5 *precision* is itself a unit of value.

The framework also makes Volume IV’s claim to progress legible at a glance: Volume III delivered partial L2 and minimal L3 across most topics; Volume IV delivered solid L0–L1 across all topics, partial L3 on the OQs, and full L4 on the infrastructure parts.

12 Path to Volume V

Volume IV reduces Volume III’s six open questions to a smaller set of named conjectures. Volume V’s principal targets are listed in table 3.

Volume V target	Source in Volume IV	What success looks like
Cosmos-internal Yoneda conjecture	Part II, reverse direction of theorem 5.2	A theorem in the Riehl–Verity cosmos $\mathcal{K}_{\text{CSS}}$ demonstrating that closure under lax pullback implies DU-Seg .
UCQ at $\aleph_\omega$	Part III, theorem 6.2	Either an explicit κ^* at which all relevant analytic-Landlands objects fit, or a counterexample showing they cannot fit at any single cardinal.
ConjectureC_rev	Part VI, deferred axiom	A constructive Henkin model, an internal completeness theorem, and a sheaf-semantic transfer principle, sufficient to discharge the reverse direction of the ZFC $\leftrightarrow$ HoTT comparison.
Merged CauchyReals module	Part V, upstream PR plan	A merged PR in the agda/cubical repository containing the seven directly-proved gaps and (where possible) the two conditional gaps, with the cubical maintainers’ approval.
Full Cubical Agda OQ1	Part I, boundary postulates B1–B4	Discharge of B1–B4 against the merged CauchyReals, producing a fully Cubical-Agda contractibility proof for $k = 1, 2, 3$ with sorry-count zero.
rk machine-check of Part II	Part II, literate rk files	Installation of the rk binary in the Volume V build environment and machine-checking of the literate files.
OQ2 and OQ6 cleanup	Volume III leftovers	OQ2 picked up via Part V’s CauchyReals; OQ6 closed via Part IV’s two-line repair.

Table 3: Volume V principal targets. Each is a named conjecture or deliverable derived from a Volume IV worker paper. The targets are *tractable* in the sense that each has a precise statement and a concrete success criterion.

12.1 Sequencing

We propose the following Volume V execution sequence:

1. **First sprint.** Merge `Cubical.HITs.CauchyReals`; discharge B1–B4 in Part I; promote OQ1’s $k = 1, 2, 3$ contractibility shadows to full Cubical Agda proofs; install the `rzk` binary and machine-check Part II’s literate files.
2. **Second sprint.** Tackle UCQ. The Rasekh–Zhu fractured condensed structure [18] is the closest existing result; the Volume V worker should examine whether the fractured structure provides a single cardinal level at which the relevant analytic-Langlands objects fit.
3. **Third sprint.** The cosmos-internal Yoneda conjecture (reverse direction of theorem 5.2). This is the hardest of the four conjectures and may itself require a Volume VI.
4. **Fourth sprint.** `ConjectureC_reverse`. This requires constructive model theory infrastructure that is not in Mathlib; a separate Volume V foundation effort may be needed.

12.2 What is unchanged

Volume IV does not remove any Volume III commitment. The HoTT-RH proposition $\text{RH} := \prod_{s:\mathbb{C}_c} \zeta(s) = 0 \rightarrow \text{IsNonTrivial}(s) \rightarrow \text{Re}(s) =_{\mathbb{R}_c} 1/2$ remains the ultimate target. The six-gate roadmap of Volume II Part VI remains structurally intact. Volume IV closes Gates 1–3 modulo `CauchyReals` and Theorem A; Volume V’s first two sprints finish those closures.

13 Risk assessment

We are explicit about Volume IV’s failure modes.

13.1 Technical risks

- **Part II’s no-go theorem may admit a more permissive formulation.** theorem 5.1 excludes proofs of **DU-Seg** that follow the GWB four-step pattern. A future proof using a different pattern — e.g., the directed-J-rule of [13] — is not excluded. The no-go is therefore relative to a strategy class, not absolute.
- **Part III’s UCQ conjecture may turn out to be unfounded.** If $\text{Sm}(\text{GL}(n, \mathbb{A}_{\mathbb{Q}}))$ in fact fits at $\aleph_1$ (contra Proposition 4.4), then Volume III’s loose claim was correct and Volume IV’s UCQ analysis is overcautious. The Volume V target remains the same (verify the cardinal level), but the narrative shifts: Volume IV’s contribution becomes a careful audit rather than a substantive correction.

- **Part V’s two conditional gaps may resist closure.** The multiplicative-inverse gap requires a Markov-style decidability assumption; if Markov’s principle is rejected (as it is in some constructive schools), the gap remains conditional in `agda/cubical` master, and the merged PR may have to ship with the conditional gaps marked. This does not block Part I, because Part I’s contractibility shadow does not require the inverse, but it does limit the scope of the merged module.
- **Part I’s $k \geq 2$ shadows may require deeper Bernoulli arithmetic than Part V supplies.** Volume III’s `decided` proofs of q_1, q_2, q_3 relied on the rationals’ computational fragment; the cubical setting may force a slightly different Bernoulli encoding that does not reduce by `decide`. If so, the Volume V cubical proofs of contractibility for $k = 2, 3$ require a re-derivation, not just a port.
- **Part VI’s Conjecture C reverse may be unprovable.** The required infrastructure (constructive Henkin model, internal completeness, sheaf-semantic transfer) has been outstanding for over a decade. If the Volume V worker concludes that no such infrastructure exists in current HoTT, then the Theorem A *forward* direction is the strongest possible result, and Conjecture C is reformulated as a *permanent* no-go rather than an open conjecture.

13.2 Methodological risks

- **The verdict-matrix arc may not survive peer review.** Our projected 12–16 valid count is conservative but not external-reviewed. A stricter reviewer might count `cache_works` as a smoke-test rather than a theorem, or the Part VI corollaries as immediate consequences of the parent theorem rather than independent verdicts. We anticipate a post-review verdict of ~ 10 valid as a defensible floor.
- **The unified framework of section 11 is descriptive, not formal.** The six-level stratification organises the volume’s content but does not formalise any new mathematics. A reviewer may ask whether the framework adds anything beyond restating the dependency diagram; our defence is that it makes the cross-cutting themes of section 10 legible as structural facts rather than rhetorical observations.
- **Cross-paper API stability is asserted, not enforced.** Part VI’s interface contract is a Lean type-class structure, so type-level changes break the build. But the prose-level narrative (e.g., “Theorem A licences transferring the q_k identities”) is not type-checked; a future change in the narrative could decouple from the Lean library without breaking anything mechanical. We rely on operational discipline (the authors are also the API consumers) for this stability.

13.3 Strategic risks

- **The volume is a half-step.** Volume IV does not prove RH, does not even close the analytic Langlands gap, does not resolve directed univalence for general Segal types. A reader expecting a full step is going to be disappointed. Our defence is that Volume IV

converts six fuzzy open questions into four sharp Volume V conjectures plus three load-bearing infrastructure items, and that this conversion is itself non-trivial work.

- **Volume V may itself be a half-step.** Even if Volume V hits all four named conjectures, RH remains gated on the analytic Langlands correspondence, the geometric Langlands $\rightarrow$ arithmetic Langlands transport, and the actual non-triviality of RH itself. A multi-volume programme is the honest expectation.

14 What this volume changes about the programme

Beyond the per-paper contributions, Volume IV makes three structural changes to the HoTT–Riemann Hypothesis programme as a whole.

14.1 From open questions to named conjectures

Volume III’s six OQs were, at their introduction, deliberately fuzzy: each was phrased to leave room for the right formalisation to be discovered. Volume IV converts the OQs into named conjectures:

- OQ1 $\rightarrow$ *The merged-CauchyReals contractibility shadow upgrades to a Cubical Agda proof.*
- OQ3 $\rightarrow$ *The cosmos-internal Yoneda conjecture is true.*
- OQ5 $\rightarrow$ *UCQ at $\aleph_\omega$ admits a positive answer.*
- Theorem A reverse $\rightarrow$ *Conjecture C admits the three named pieces of constructive infrastructure.*
- OQ2, OQ6 $\rightarrow$ closed via Volume IV’s infrastructure.

The shift is from “what is the right formalisation?” to “does this specific named statement hold?” This is a strict refinement.

14.2 From single-paper deliverables to API stacks

Volume III published seven papers each with its own Lean stub and its own self-contained statements; cross-paper sharing happened in prose, not in code. Volume IV introduces a deliberate stack: Part IV is the toolchain; Parts V–VI are the foundation libraries; Parts I, III consume those libraries by name. The stack is enforceable at build-time: a downstream change that breaks the API breaks the build.

This pattern is standard in software engineering but uncommon in mathematics. The argument for adopting it is the same as in software: shared dependencies amortise, and an explicit interface is machine-checkable in a way that prose is not.

14.3 From verdict counts to verdict structure

Volume III’s verdict matrix was a flat 3/0/29. Volume IV introduces the four-axis matrix proposed in section 10 (valid / invalid / statement-incomplete / toolchain-incomplete) and the six-level stratification of section 11 as an auxiliary view. Both add structure to what was a flat count.

The shift matters because flat counts conflate qualitatively different kinds of progress. A future reader of Volume III’s 3/0/29 might conclude that 90% of the matrix was open mathematical problems; a four-axis reading reveals that perhaps 60% was intentional scaffolding and 25% was build-broken, leaving only about 15% as genuine open mathematics. Volume IV’s matrix should be read with the same care.

15 Comparison with related programmes

Volume IV’s approach — combining HoTT-native formulation with Lean shadowing, infrastructure investment, and explicit no-go theorems — has parallels in several ongoing formal-mathematics programmes.

15.1 Mathlib and the broader Lean ecosystem

Mathlib’s growth strategy has been to maintain a fast-moving master branch and ask downstream consumers to track or pin. Volume IV’s Part IV adopts the pinning posture explicitly. The trade-off — pinning loses access to upstream improvements, but gains reproducibility — is the right one for a multi-volume programme in which each volume must be re-runnable years after publication.

15.2 The Liquid Tensor Experiment and Polynomial Functors

The Liquid Tensor Experiment [25] formalised Scholze’s main theorem of liquid vector spaces in Lean; Polynomial Functors [26] formalised a substantial ∞ -categorical fragment. Both demonstrate that a sufficiently disciplined Lean project can formalise content beyond Mathlib’s current scope. Volume IV’s Part III takes a more cautious posture: the $(\infty, 1)$ -topos $\mathcal{E}_{\text{cond}}$ is not formalised in Lean; only the theorem statements are, with proofs marked `sorry`. The reason is that condensed mathematics in Lean is itself a substantial effort (Topiary, etc.) that is out of Volume IV’s scope. Volume V may revisit this.

15.3 The agda/cubical project

The `agda/cubical` repository has a high merge bar; Booij and Mörtberg’s `CauchyReals` module sat unmerged for an extended period because the full module exceeds the maintainers’ sustainable review load. Volume IV’s Part V re-engages the merge effort, with the expectation that the seven directly-proved gaps form a defensible sub-PR even if the two conditional gaps are deferred. The operational lesson is that upstream infrastructure investment requires explicit allocation of worker time, not just citation in a “future work” section.

15.4 The Riehl–Verity cosmos and `rzk`

The `rzk` proof assistant [15, 16] implements a variant of STT/TTT and has been used for the sHoTT and Yoneda formalisations. Volume IV’s Part II writes its formalisation in `rzk` syntax but does not machine-check it (the binary was not installed in the Volume IV build environment). Volume V should remediate. The broader observation is that the synthetic- ∞ -category formalism has matured to the point where non-trivial directed-univalence content can be expressed in a real proof assistant; Part II’s no-go theorem benefits from this maturity in that the proof structure is recognisable to anyone familiar with sHoTT.

16 Discussion

16.1 What Volume IV is

Volume IV is a half-step. It converts Volume III’s collection of formulated open questions into a smaller collection of tractable named conjectures, plus standing infrastructure for the rest of the programme. The conversion is non-trivial but not heroic; the Volume V deliverables remain hard.

The volume’s principal honest finding is that $\mathcal{E}_{\text{cond}}$ *is* elementary in the universe-stratified sense, with UCQ at $\aleph_\omega$ as the tight remaining gap. This sharpening is the volume’s most consequential mathematical contribution: it identifies the cardinal coherence, not the local axiom check, as the actual obstruction to a HoTT-internal analytic Langlands theory.

The volume’s second principal contribution is operational: the toolchain-pinning pattern, the citation-API pattern, and the four-axis verdict matrix are deliverables that any subsequent volume of any HoTT-formalisation programme can adopt directly. These are infrastructure, in the broadest sense.

16.2 What Volume IV is not

Volume IV is not a proof of RH. It is not a proof of directed univalence for general Segal types; in fact theorem 5.1 shows that one specific class of strategies cannot work, and the Riehl–Verity reduction is open in the reverse direction. It is not a complete elementary $(\infty, 1)$ -topos certificate for $\mathcal{E}_{\text{cond}}$; it is an audit that succeeds at five axioms in the stratified sense but defers the cardinal coherence to Volume V. It is not a constructive Henkin model for ConjectureC_reverse; it documents what such a model would have to look like and exposes the gap as an axiom.

The volume is honest about all of this. The Risk Assessment of section 13 flags six technical risks and three methodological risks. The unified framework of section 11 explicitly distinguishes proofs (L3) from compositions (L4) from named conjectures (L5).

16.3 The synthesis as a deliverable

The present synthesis paper is itself the L4-deliverable for the volume as a whole. Each Volume IV worker paper is a contribution at some level (Part IV $\rightarrow$ L0; Parts V, VI $\rightarrow$ L1; Parts I, II, III $\rightarrow$ L2–L3); the synthesis is the L4 layer that makes the parts compose.

The Volume IV abstract reads as “six topics that collectively advance the programme”; the synthesis reads as “the six topics fit together in this way and reduce the programme to these four named conjectures.” Without the synthesis, the volume is six papers; with it, the volume is a stage.

17 Conclusion

Volume IV of the HoTT–Riemann Hypothesis programme delivers six worker papers that, in combination, advance the programme from Volume III’s collection of formulated open questions to a smaller collection of named conjectures plus standing infrastructure for the rest of the programme.

The principal Volume IV contributions:

1. **Part IV (infra-mathlib-cache):** Mathlib pin, pre-warmed cache, CI workflow. `lake build` succeeds in 4.8 seconds with 8020 cached oleans. Unblocks all four Lean-based topics.
2. **Part V (infra-cauchy-reals):** Seven gaps closed by direct proof, two reduced to conditional statements. Worked $\sqrt{2}$ example. Lean shadow with three sorry-free theorems. Upstream PR plan to `agda/cubical`.
3. **Part VI (infra-comparison-lemmas):** Theorem A forward and reverse, Theorem B forward, six corollaries, all sorry-free; `ConjectureC_reverse` as explicit deferred axiom. `lake build` succeeds with 3,307 jobs.
4. **Part I (oq1-cubical-zeta):** Cubical Agda starter with B1–B4 boundary postulates gated on `CauchyReals`. Lean shadow contractibility for $k = 1$, sorry-count zero. Constructive Bernoulli library.
5. **Part II (oq3-directed-univalence):** No-go theorem (theorem 5.1); half-reduction theorem (theorem 5.2); rep-theoretic corollary; reverse direction identified with cosmos-internal Yoneda conjecture.
6. **Part III (oq5-langlands-topos):** Verdict *universe-stratified elementary* for $\mathcal{E}_{\text{cond}}$; UCQ at $\aleph_\omega$ as the tight remaining gap; HoTT-internal automorphic representations conditional on UCQ; liquid/solid obstruction generalisation; conditional OQ3+OQ5 composition.

The five cross-cutting themes (toolchain pinning, the `CauchyReals` chain, the Theorem A citation API, the verdict-matrix arc, no-go as contribution) and the unified six-level framework (toolchain $\rightarrow$ foundation $\rightarrow$ statements $\rightarrow$ proofs $\rightarrow$ composition $\rightarrow$ Volume V interface) organise the per-paper contributions into a single coherent deliverable.

The path to Volume V is explicit. The seven named targets are: the cosmos-internal Yoneda conjecture; UCQ at $\aleph_\omega$; `ConjectureC_rev`; the merged `CauchyReals` module; the full Cubical Agda formalisation of OQ1; the `rzk` machine-check of Part II; and the OQ2/OQ6 cleanup. Each is a named deliverable with a precise success criterion.

The honest summary of the volume is the one given in the abstract: the volume is a half-step, deliberately so, that converts six fuzzy open questions into four sharp conjectures plus three load-bearing infrastructure items. The programme remains gated on the analytic Langlands correspondence and on RH itself. But the gates are sharper, the infrastructure stands, and the next volume’s targets are tractable in a way that the previous volume’s targets were not.

Acknowledgements

The author thanks the YonedaAI Research Collective for sustained support of the HoTT–Riemann Hypothesis programme. Phase 3 review of the six worker papers was conducted by Gemini 3.1 Pro; LaTeX formatting was reviewed by Codex (OpenAI). All errors are the author’s.

References

- [1] M. Long, *OQ1 — Full Cubical Agda Formalisation of Coinductive $\zeta(2k)$: From Formulation to Formal Proof*, Volume IV Part I, *The YonedaAI Collaboration*, 2026.
- [2] M. Long, *OQ3 — Directed Univalence for Non-Discrete Types: A No-Go Theorem and a Half-Reduction*, Volume IV Part II, *The YonedaAI Collaboration*, 2026.
- [3] M. Long, *Identifying the $(\infty, 1)$ -Topos for Analytic Langlands: A Rasekh-Axiom Audit of $\mathcal{E}_{\text{cond}}$ with Cardinal-Stratification Caveats*, Volume IV Part III, *The YonedaAI Collaboration*, 2026.
- [4] M. Long, *Cross-Cutting Infrastructure I: A Pre-Warmed Mathlib Build Cache for Volume IV of the HoTT–Riemann Programme*, Volume IV Part IV, *The YonedaAI Collaboration*, 2026.
- [5] M. Long, *Cross-Cutting Infrastructure II: Completing `Cubical.HITs.CauchyReals`*, Volume IV Part V, *The YonedaAI Collaboration*, 2026.
- [6] M. Long, *Cross-Cutting Infrastructure III: The Comparison Lemma Library: From Theorem A to a Stable Citation API*, Volume IV Part VI, *The YonedaAI Collaboration*, 2026.
- [7] *Volume III Lean Verdict Matrix*, `reviews/lean-verdicts.md`, *hott-riemann-hypothesis* repository, 2026.
- [8] *OQ1 Lean Verdict*, `reviews/oq1-coalgebraic-zeta-2k-lean-verdict.md`, *hott-riemann-hypothesis* repository, 2026.
- [9] *OQ5 Lean Verdict*, `reviews/oq5-infinity-topos-analytic-langlands-lean-verdict.md`, *hott-riemann-hypothesis* repository, 2026.

- [10] M. Long, *Coinductive Witnesses for $\zeta(2k)$: A Streaming Presentation of the Euler–Bernoulli Identity in HoTT*, Volume III OQ1, *The YonedaAI Collaboration*, 2026.
- [11] D. Gratzer, J. Weinberger, U. Buchholtz, *Directed Univalence in Simplicial Homotopy Type Theory*, arXiv:2407.09146, 2024.
- [12] U. Buchholtz, D. Gratzer, *Cocompleteness in Simplicial Homotopy Type Theory*, TYPES 2025.
- [13] T. Altenkirch, J. Neumann, *Synthetic 1-Categories in Directed Type Theory*, arXiv:2410.19520, 2024.
- [14] E. Riehl, D. Verity, *Elements of ∞ -Category Theory*, Cambridge University Press, 2022.
- [15] E. Riehl, M. Shulman, *A Type Theory for Synthetic ∞ -Categories*, arXiv:1705.07442, 2017.
- [16] N. Kudasov, J. Weinberger, J. Neumann, *rzk: a proof assistant for synthetic ∞ -categories*, project documentation, [rzk-lang.github.io](https://github.com/rzk-lang/rzk-lang), 2025.
- [17] N. Rasekh, *A Theory of Elementary Higher Toposes*, arXiv:1805.03805, 2018.
- [18] N. Rasekh, B. Zhu, *Fractured Structures in Condensed Mathematics*, arXiv:2603.09618, 2026.
- [19] D. Clausen, P. Scholze, *Lectures on Condensed Mathematics*, Bonn lecture notes, 2019.
- [20] D. Clausen, P. Scholze, *Lectures on Analytic Stacks*, ongoing, 2024.
- [21] D. Gaitsgory, S. Raskin, et al., *Proof of the Geometric Langlands Conjecture, I–V*, arXiv:2405.03599 and four follow-ons, 2024–2025.
- [22] A. Booij, *Analysis in Univalent Type Theory*, PhD thesis, University of Birmingham, 2020.
- [23] A. Mörtberg, A. Vezzosi, *Cubical Agda: A Dependently Typed Programming Language with Univalence and Higher Inductive Types*, JFP 31, 2019.
- [24] S. Doré, E. Cavallo, A. Mörtberg, *Automating Boundary Filling in Cubical Agda*, FSCD 2024.
- [25] D. Castelvechi, *Mathematicians welcome computer-assisted proof in ‘grand unification’ theory*, Nature 595, 18–19, 2021.
- [26] J. Avigad et al., *Polynomial Functors: A Mathematical Theory of Interaction*, ongoing Lean formalisation, 2024.
- [27] J.-P. Serre, *Linear Representations of Finite Groups*, Springer GTM 42, 1977.
- [28] R. Diaconescu, *Axiom of choice and complementation*, Proc. AMS 51, 175–178, 1975.