Unfaithfulness in the Quasi-Riemann formalization

HumanLean Team

November 2026

What counts as unfaithful?

A formal proof establishes the claim written in code. Faithfulness asks whether that claim and its proof preserve the mathematics of the source text. We assess two dimensions separately:

- **Definition and statement faithfulness:** the formalization preserves the intended objects, assumptions and conclusions. A changed construction or weaker conclusion can alter what the result establishes.
- **Proof faithfulness:** the formalization preserves the substantive argument of the source proof. Routine changes in notation or implementation do not count. Replacing the central argument with another method counts as a proof difference, even when both methods prove the same theorem correctly.

Findings concern the numbered result or its own argument. Differences inherited solely from dependencies are excluded. Direct reuse of existing Mathlib declarations and transparent wrappers or delegations to them are excluded.

Findings

In the formalization of *The Quasi-Riemann Hypothesis*, we identify substantive differences in **9 numbered results**:

- **4** results have statement or mathematical-object mismatches.
- **8** results have substantive proof replacements.

Three results have both kinds of finding. The semantic findings concern what the formalization expresses. The proof findings concern how it establishes its conclusion.

The findings below follow the sections of the paper. Each entry explains the difference and links the original argument and formal code needed to check it.

Section 4. Arithmetic and analytic preliminaries

Lemma 4.1 — Ray characters and the Kummer extension

Semantic. For a nonzero Eisenstein integer a, the source constructs the extension $F(a^{1/6})/F$, where $F = \mathbb{Q}(\sqrt{-3})$. It asserts that this extension is abelian, unramified outside $6a$, and identifies arithmetic Frobenius with the sixth-power residue symbol. The ray character is a consequence of those field-theoretic assertions. The mapped formal construction produces a finite-residue character by proving periodicity and descending to a quotient. This supplies the ray-character consequence without supplying the lemma's extension, ramification and Frobenius clauses.

Proof. The source uses the action on a sixth root, the discriminant and Artin reciprocity. The formal construction uses explicit reciprocity, supplementary laws and finite-residue descent. This arithmetic route agrees with a later ray-character argument in the paper, which is credited. It changes the original lemma's field-theoretic route. The source identifies characters only on ideals coprime to $6a$, retains zeros at collisions with a, and asserts no conductor-exponent bound. The formal periodic-row construction also retains coprimality masks.

Paper pp. 11–12 · Source statement and proof · Later reciprocity argument · RowClosure.lean:11–79 · PeriodicRows.lean:117–166 · RowClosure.lean:216–234

Lemma 4.6 — Gaussian annular decomposition

Proof. The source cuts the logarithmic Gaussian into compact annular profiles. Repeated real differentiation gives a polynomial times the same Gaussian, yielding summability after every fixed exponential weight and tails smaller than every prescribed power of the large scale. The mapped formal proof extends the Gaussian to a complex analytic function and applies Cauchy derivative estimates on disks. It then uses arbitrary power decay and weighted geometric dyadic sums to control the tails.

Both routes exploit the Gaussian's rapid decay. The substantive difference is the all-orders derivative and tail argument, beyond a choice of annular indexing. The formal implementation's Gaussian profile and derivative estimates are credited as alternate implementations of the required analytic bounds.

Paper pp. 19–20 · Source statement and proof · GaussianFlow.lean:35–82 · GaussianTailBudget.lean:18–74

Section 5. Completed cubic reflection and unmarked row energy

Lemma 5.4 — Generic lattice tails versus a structured kernel

Semantic. The source permits any kernel K with a single real-order bound $|K(x)| \le C_A x^{-A}$ for $x \ge 1$, where $A > 1$. For either specified Eisenstein lattice, $a > 0$ and $U \ge 1$, it bounds the absolute tail by a constant times $C_A(1 + a^{-1})U^{1-A}$. The mapped formal estimate concerns the structured reflection kernel `paperKernel`. For an integer tail order N, it requests decay of order $N+2$ and obtains reciprocal-square-scale dependence, explicitly $(\text{scale}^2\rho^6)^{-1}$.

The source controls a whole class of kernels using one decay assumption and specified uniform dependence on the scale. Availability of arbitrarily high decay for this particular smooth kernel supports its application, while leaving the generic single-order assertion uncovered by the mapped estimate.

Proof. The source counts lattice points in successive shells and sums two geometric series. The formal proof extracts the tail power and sums products of inverse ideal norms together with a geometric factor.

This proof difference accompanies the changed kernel contract. The elementary summation variant alone is not assigned independent major severity.

Paper p. 36 · Source statement and proof · RawTail.lean:11–85 · UniformKernelBounds.lean:315–362

Lemma 5.5 — Quadratic reduction by total products

Proof. The source first proves the coefficient-sensitive estimate (5.27). It factors $b = gt^2$, separates $c = (n, g)$, expands the collision mask by Möbius inversion, and applies a rescaled quadratic large sieve. Summing those estimates gives the bounded-coefficient consequence (5.28).

The mapped formal proof of (5.28) first groups all pairs (n, b) by their total ideal product. It bounds the aggregated coefficient by the number of divisors, freezes the product's square part, and combines the resulting sieve bounds by weighted Cauchy. This changes the central reduction: it estimates the aggregated products directly, bypassing the refined collision-mask expansion used to derive (5.27).

The bounded-coefficient endpoint is retained. The contrast concerns its proof, including how collisions are handled. The formal square-part calculation keeps a zero mask in its exact identity before bounding its norm by one.

Paper pp. 37–38 · Source statement and proof · UniformTailBounds.lean:77–120 · ProductColumns.lean:77–121 · ProductColumns.lean:546–711

Section 6. The base probe and its balanced low estimate

Lemma 6.1 — Geometric large sieve versus arithmetic correlations

Proof. The source proves a planar additive large sieve for arbitrary δ-separated points on $\mathbb{C}/\mathcal{O}$. Its bound has size $Q + \delta^{-2}$: frequency-ball size and point separation both control the energy. The proof constructs a bandlimited majorant, applies Poisson summation, and bounds nearby lifted points by packing disjoint disks.

The mapped low-side estimate instead expands arithmetic Gauss correlations, separates exceptional and nonexceptional frequencies, and bounds divisor and ideal sums. The replacement changes the mechanism used for the Gram bound from separated-point geometry to the arithmetic structure of these rows.

The library also contains Huxley finite primal and dual large-sieve estimates for reduced rational frequencies. Those are credited. This finding concerns the route invoked for the low-side bound and carries only the proof dimension.

Paper pp. 47–48 · Source geometric proof · GramBound.lean:13–59 · GramSourceEstimate.lean:12–41 · GramIdealMass.lean:30–107 · HuxleyFiniteDual.lean:56–110

Proposition 6.3 — Construction of the low bound

Proof. The source estimates the balanced probe at $X = Y = Z^{1/2}$ by $Z^{1/4+\epsilon}$. It applies the completed Gaussian moment directly, pairs it with the additive norm using Cauchy–Schwarz, and integrates the Mellin weight. The original row zero masks remain in that pairing.

The mapped construction estimates a compensated probe at $X = Z^{17/48}$ and $Y = Z^{23/48}$, with low exponent 3/16 plus arbitrarily small loss after normalization. It reconstructs the probe by summing

over subsets of auxiliary prime slots and by combining central and remote Gaussian dyads. The substantive local difference is this reconstruction and low-bound assembly, beyond the changed Gram or reflected-energy inputs.

The Cauchy–Schwarz and Mellin transports, together with exact mask transports, remain faithful components. The compensated geometry also appears explicitly in the paper's second part. The finding compares this construction with the balanced argument of the numbered proposition.

Paper p. 49 · Source balanced argument · Compensated geometry in Part II · LowCentralTuple.lean:18–69 · LowCommonBound.lean:23–75 · LowNormalized.lean:13–36

Section 9. A sextic-sieve row count

Proposition 9.2 — A later amplification route for counting rows

Proof. The source uses the inverse witness at fixed $t = 1$. It factors each sixth-power-free row into a squarefree part and a powerful part, whose positive prime valuations are at least two. It applies the squarefree sextic sieve on each powerful-part stratum and optimizes the resulting exponent envelope. The floor bin is handled by elementary counting.

The mapped operative count uses sixth-power amplification and combines inverse, plain and prime-marked witness counts, optimizing with an amplitude parameter. Amplification maps a sixth-power-free row to its product with a sixth power. Unique factorization supplies the injectivity needed to count these enlarged rows. The formal prime-deletion identities retain the associated zero masks.

This is a recognizable implementation of the paper's later amplification and refined-count arguments. It changes the central argument of Proposition 9.2, which uses only the inverse witness and explicitly needs no plain-witness product. Borrowing a correct later argument still produces a proof-route difference for this proposition.

Paper pp. 67–70 · Source row-count proof · Later sixth-power amplification · Later refined counts · InverseAmplificationRows.lean:10–74 · InverseAmplificationFamily.lean:31–61 · DetectorCountFromMoments.lean:20–90

Section 11. The 11/12 conclusion

Proposition 11.2 — Balanced and compensated high estimates

Semantic. The source estimates the balanced normalized probe $J_{\mathrm{I},\eta}$ against its residue signal. Its signal exponent is $C_{\mathrm{I}}(s) = s - 2/3$, its scales are $X = Y = Z^{1/2}$, and its error saves $1/4800$ in the exponent under the positive-gap supposition of Part I. The mapped estimate uses a compensated prime-slot probe at $X = Z^{17/48}$, $Y = Z^{23/48}$, with signal shift $-11/16$ and a positive saving selected through the formal parameter budgets.

A probe is a smoothed arithmetic sum designed to convert zeros into an asymptotic signal. Changing its scales and compensating prime factors changes the intermediate assertion being estimated. The compensated probe is explicitly distinguished from the balanced probe in Part II of the source paper.

Proof. The source assembles the balanced principal residue, intermediate-row and outer-range estimates, then closes the height and loss budgets. Its base principal correction has residue error identically zero. The formal high-bound assembly uses the compensated contour geometry and a prime-slot principal-residue comparison with additional error terms. This local analytic comparison changes the high-bound argument independently of the different row-count input.

Paper pp. 83–84 · Source high estimate · Balanced probe and shift · Base principal correction · Part II probe and shift · FinalAssemblySource.lean:15–25 · FinalAssemblyHigh.lean:98–133 · FinalAssemblyFixedHigh.lean:14–45 · NormalizedTransport.lean:72–119

Section 17. The inverse moment with prime factors

Corollary 17.4 — Abstract propagation versus canonical energies

Semantic. The source gives an abstract propagation principle on a finite acyclic graph. Vertices carry arbitrary nonnegative quantities, subject to the specified terminal and recursive bounds. Edges may carry indexed kernels, nonnegative measures and real child powers, with uniform profile and height control. After normalization by the desired exponents, the corollary yields finite seminorm and height orders independent of an external cutoff.

The mapped formal endpoint bounds canonical Eisenstein row energies. Its contract fixes the arithmetic coefficient class, prime-pool weights, rank steps, row caps and loss schedules. It supplies a concrete arithmetic application of propagation, leaving the general vertex-and-edge contract uncovered by this endpoint. In intuitive terms, a theorem about how a whole class of recursive estimates can be combined has been specialized to one structured family of energies.

Backward choice of orders and finite-rank induction preserve the source's propagation mechanism and receive no proof finding. The source also requires a coefficient measure common to the relevant rows and normalizes genuine outer labels so their mass is counted once. These conventions are part of the abstract contract being compared.

Paper pp. 118–119 · Abstract graph hypotheses · Source corollary and proof · Common measure and label normalization · CanonicalRankExistenceData.lean:15–49 · CanonicalRankExistenceInduction.lean:23–54 · CanonicalRankExistence.lean:21–38

Sources and verification

The comparison uses the paper and formalization at revision `4eaee87`. It lists retained findings from a screening of 63 numbered results and covers the original two-part argument, including Part I as a dependency of Part II. The application in Corollary 1.2, alternative manuscripts and separate Siegel results are outside the scope. Evidence links identify exact paper pages and code lines. This is a source-level comparison, without a new Lean compilation or kernel check.