The Invariant Structure of the Modular Surface

The Tetrahedral Framework, the Selberg Spectrum, and the Boundary at the Riemann Zeros

QRiemannian Collaboration — Andri Sigurgeirsson Vidalin & Claude

QRiemannian Research — June 2026 (v2 register edition, July 2026)

The Invariant Structure of the Modular Surface: the Tetrahedral Framework, the Selberg Spectrum, and the Boundary at the Riemann Zeros

Abstract

The QRiemannian “tetrahedral framework” is here examined claim-by-claim against established mathematics, and characterized for what it provably is. Read coldly, its distinctive content resolves to the invariant structure of a single object: the modular surface X(1) = SL(2,$\mathbb{Z}$)\$\mathbb{H}$ and the real-quadratic field $\mathbb{Q}$(√5). The eigenvalue trinity (π, √5, φ) and the coupling g_c are the finite invariants of X(1) — area, discriminant, systole, and their forced combination [CV]/[D]; the field $\mathbb{Q}$(√5) is a richly-chosen, canonically-located marked point of that surface, not a freely selected one [D]. The framework’s commutator algebra genuinely reaches the Selberg/Maass spectrum: the sl(2,$\mathbb{R}$) Casimir is the hyperbolic Laplacian with spectrum $\lambda = ¼ + r^2$, and the framework’s transfer operator is the Mayer operator detecting Selberg-zeta — not Riemann — zeros [CV]. But it forces nothing there beyond that single marked-point fact [D]: its golden content sits in the length-spectrum channel and is structurally barred from the channel where the Riemann zeros live. The gap between this sound description and the Riemann Hypothesis is exactly RH itself.

We state the firewall up front. This paper claims none of the following: a proof of the Riemann Hypothesis; a Hilbert–Pólya operator; that the framework derives the Riemann zeros; or that g_c is a physical coupling. The contribution is a rigorous characterization together with a precisely-named boundary. A precisely-named gap is a result.

1. Introduction

The object of this paper is a framework that names three numbers — π, $\sqrt{5}$, and the golden ratio φ — and a coupling g_c, and organizes them into a structure it calls tetrahedral. We characterize what that structure mathematically is. The answer is specific and checkable: examined claim-by-claim against established mathematics, the framework's distinctive content resolves to the invariant structure of the modular surface $X(1) = SL(2,\mathbb{Z})\backslash\mathbb{H}$ and the real-quadratic field $\mathbb{Q}(\sqrt{5})$.

Three claims carry that characterization, each at its own grade. First, the trinity (π, $\sqrt{5}$, φ) and the coupling g_c are the finite invariants of X(1) — π from the area, $\sqrt{5}$ from the discriminant, φ from the systole, g_c as their forced combination — and not numerical coincidences [CV]/[D]. Second, $\mathbb{Q}(\sqrt{5})$ is a marked point of that surface: a forced home (the universality of X(1) as the moduli of elliptic curves) together with a chosen spot (the shortest geodesic). It is richly and canonically located, not arbitrarily selected — but the locating is "marked," not "derived from a selection principle," and we hold it at that grade [D]. Third, the framework's $sl(2,\mathbb{R})$ commutator algebra genuinely reaches the Selberg/Maass spectrum: the Casimir element is the hyperbolic Laplacian, with spectrum $\lambda = ¼ + r^2$, and the framework's transfer operator coincides with the Mayer transfer operator, which detects the zeros of the Selberg zeta function [CV].

The reach is real. It is also bounded, and the boundary is the point of the paper. The golden content — the $\sqrt{5}$ and φ that distinguish this framework from generic modular-surface bookkeeping — is confined to the length-spectrum channel of X(1) and is structurally barred from the channel in which the Riemann zeros appear (the scattering block, built from ζ, Γ, and π, with no $\sqrt{5}$ entering). The single arithmetic route that would carry $\sqrt{5}$ into a zeta function changes the surface itself, away from X(1). The framework therefore reaches the spectrum without reaching the zeros — and the two are not the same reach. The distance between this sound description and the Riemann Hypothesis is not a missing lemma the framework nearly supplies; it is RH itself, un-crossed. We hold the full characterization at [D] — a characterization consistent across independent analyses (strong corroboration, not a proof).

Scope and firewall. The contribution is a characterization with a precisely-drawn boundary, and the boundary is part of the claim, stated at the outset. This paper does not claim a proof of the Riemann Hypothesis; it does not construct a Hilbert–Pólya operator; it does not claim that the framework derives the Riemann zeros; and it does not claim that g_c is a physical coupling. The framework's broader physical interpretation is preserved elsewhere as an interpretation [I]; here we characterize only the mathematics it would rest on. A distinct layer — the claims that the structure forces the zeros, fixes external physical constants, or supplies independent confirmations — was examined, is separately graded FIRED, and is excluded here. The boundary is drawn in both directions: the object-layer is sound, its internal arithmetic closed and independently recomputed, and the trinity invariants established or derived-within-framework [CV]/[D]. The paper's strength is precision about where each claim stands — neither a proof nor a dismissal.

Method. The characterization is not asserted; it is the convergent outcome of a systematic stress-test. Six structurally-independent investigations — of the marked-point anchor, of whether any functional selects it, of the framework's operator structure against a candidate obstruction, of the coupling g_c, of the trace-formula content, and of the Selberg/Maass landing — each ran to a graded conclusion, and all six landed on one shape: a genuinely rich, internally sound arithmetic-geometric object that describes the modular surface and reaches the Selberg/Maass spectrum, but does not force the Riemann zeros. The agreement across independent investigations strengthens confidence in the shape of the finding; it does not raise its grade to proof. We say so wherever it matters.

The six investigations, in plain terms. The six independent lines of inquiry, each approaching the framework from a different direction, were: (1) the marked-point anchor — an examination of the arithmetic that ties the framework's location to $\mathbb{Q}(\sqrt{5})$ (the smallest-trace hyperbolic class of X(1), and the identity $3^2 - 4 = 5$), asking how forced that anchoring actually is; (2) the search for a selecting functional — a constructive attempt to find any principle, variational or otherwise, that would select the systole as an output rather than merely read it as an input; (3) the operator-structure confrontation — a test of the framework's commutator and transfer-operator apparatus against a candidate spectral obstruction, to see whether it reaches a self-adjoint operator for the Riemann zeros or only the Selberg/Maass spectrum; (4) the coupling g_c — a grading of whether $g_c = 1/(\sqrt{5}\cdot\pi)$ carries genuine derived content or is a bare product of two independent invariants; (5) the trace-formula content — a check of the framework's reading against the standard Selberg-trace-formula structure of X(1) (the area, the elliptic fixed points, the cusp); and (6) the Selberg/Maass landing — a verification of where the framework's algebra lands on the automorphic spectral theory of the surface. Each ran to a graded conclusion by a different route, and all six arrived at the same shape.

Grading convention. Every load-bearing claim in this paper carries an explicit grade, used uniformly throughout:

[CV] — established / conventional: a result of established mathematics, cited or directly checkable.

[E] — exact identity: an exact closed-form equality, true by direct computation; often written together with [CV] or [D] to mark that an established or derived claim is also exact.

[D] — derived within this work: follows by sound reasoning from the structure under study, but is not an external theorem.

[C] — structural proposal made in this work: a proposed identification or correspondence, offered as structure, not established. (This paper's claims fall under the other grades; [C] is included so the key is identical to that of the companion paper.)

[I] — interpretation: a reading of the structure, not a derivation from it.

FIRED — examined and excluded: tested against its falsifier, found not to hold as stated, and excluded from the paper's claims.

The grades are not decoration. A characterization is only as honest as the boundary it draws, and the grades are how the boundary is drawn — claim by claim, including the claims that did not survive.

2. The modular surface and its marked point

The framework's home surface is the modular surface

$X(1) = SL(2,\mathbb{Z})\backslash\mathbb{H}$.

This object is not chosen for the framework's convenience. It is, independently of any framework consideration, the universal modular surface: the moduli space of elliptic curves, the quotient of the upper half-plane $\mathbb{H}$ by the full modular group $SL(2,\mathbb{Z})$. Whatever distinctive content the framework carries lives on this surface, and the surface's privilege is established mathematics, not a framework posit [CV]/[E]. We take this as the bedrock of the characterization, and we are careful below to keep it separate from the question of where on the surface the framework sits.

2.1 The forced home

That X(1) is the natural home is a matter of standard arithmetic geometry. As the moduli of elliptic curves over $\mathbb{C}$, and as the quotient by the full $SL(2,\mathbb{Z})$ — not a congruence subgroup, not a Fuchsian group chosen to taste — X(1) is fixed before the framework speaks. The framework's further restriction to degree-2 arithmetic and to closed geodesics is earned geometrically: closed geodesics on X(1) correspond to conjugacy classes of hyperbolic elements of $SL(2,\mathbb{Z})$, and these correspond in turn to (classes of) binary quadratic forms and to periodic continued fractions — the classical class → geodesic → form dictionary (Sarnak; Series). Quadratic surds are exactly what give closed geodesics, so the degree-2 locus is where the closed-geodesic structure lives; the restriction is a consequence of the geometry.

We state the home claim at its grade and no higher. X(1) is the forced home — the universal modular surface, fixed independently of the framework [CV]/[E]. This is the level surface on which everything that follows is read. It is emphatically not a selection principle for $\mathbb{Q}(\sqrt{5})$: the universality of the surface says nothing, by itself, about which point of the surface the framework occupies. Conflating the two — letting the genuine privilege of the home stand in for a privilege of the spot — is the error this section is built to avoid.

2.2 The chosen spot: the shortest geodesic

Within the forced home, the framework sits at one distinguished location: the shortest closed geodesic — the systole of X(1).

The systolic geodesic corresponds to the hyperbolic conjugacy class of smallest trace. In $SL(2,\mathbb{Z})$ a hyperbolic element has integer trace with $|t| > 2$, so the smallest available trace is $t = 3$. A representative of this class is

$F^2 = [[2, 1], [1, 1]]$, $\mathrm{tr}\, F^2 = 3$, $\det F^2 = 1$, an element of $SL(2,\mathbb{Z})$ with eigenvalues φ^2 and φ^{-2}, where $\varphi = (1+\sqrt{5})/2$ is the golden ratio. (The matrix is the square of the form whose entries are consecutive Fibonacci numbers; we write the trace-3, determinant-1 element directly, since it is the systolic class that carries the geometry.) Its two fixed points on the boundary $\partial\mathbb{H} = \mathbb{R}\cup\{\infty\}$ are φ and $-1/\varphi$.

The translation length of a hyperbolic element with eigenvalue λ is $\ell = 2 \log \lambda$. With $\lambda = \varphi^2$ this gives the systole

$\text{systole}(X(1)) = 2 \log(\varphi^2) = 4 \log \varphi \approx 1.9248.$

The length and the norm are consistent at the same power of φ: the systolic class has $N(F^2) = \varphi^4$, matching $\ell = 4 \log \varphi$ through $\ell = 2 \log \lambda$ with $\lambda = \varphi^2$. The arithmetic of the systole is the arithmetic of the golden unit: $N(\varphi) = \varphi\bar{\varphi} = -1$ and $N(\varphi^2) = +1$, so $\varphi^2 = (3+\sqrt{5})/2$ is the fundamental automorph. The factor of 4 in $4 \log \varphi$ resolves into two distinct factors of 2: one from the universal length formula $\ell = 2 \log \lambda$ (which holds for every hyperbolic element, independent of the norm), and one from the squaring $\lambda = \varphi \to \varphi^2$ that lifts the norm-(−1) unit — whose representative [[1,1],[1,0]] has det = −1 — to the determinant-1 systolic element F^2.

This is the spot. It is reached by one move repeated at every rung — take the shortest geodesic, equivalently take the smallest — and that move is a free extremization: it is a framework commitment, recurring without an external enforcer that makes the smallness consequential. The location is genuine and exact; what it is not is the output of a principle that selects it.

2.3 The marked-point verdict

Putting the two together: the framework has a forced home and a chosen spot. The honest word for $\mathbb{Q}(\sqrt{5})$'s position on X(1) is marked, not selected [D].

The distinction is precise. A selected point would be the distinguished output of some principle — a functional whose minimization, or a dynamics whose attractor, lands on the systole and would land elsewhere on a different surface, predicting beyond the framework's own training. A marked point is one the framework simply sits at: the location is real, canonical, and richly structured, but it is fixed by the choice "shortest," not derived from a selecting principle. The framework is marked at $\mathbb{Q}(\sqrt{5})$. The arbitrariness in that position is compressed enormously — from "why this field, among all number fields?" down to "why the shortest geodesic of the forced surface?" — but it is not compressed to zero. The residue is the single word shortest.

We hold this at [D] with a reopener, and the reopener is load-bearing for the grade. The marked-point verdict is not a theorem that no selecting principle can exist. It is the finding that, across the routes the investigation could construct, no functional built from established-mathematics ingredients together with framework-implied ingredients selects the systole as an output rather than reading it as an input; every candidate reads the geometry it is handed. A future selecting dynamics — variational or non-variational — that cleared this ledger would reopen the stronger "selected" reading. None was found across the investigation's routes. The grade is therefore [D], not an established impossibility: the marked point stands where the analysis can see, and the door back to "selected" is held honestly open. The detailed structure of that reopener — the precise gap, and why importing the missing structure would amount to manufacturing the conclusion — belongs to the discussion (§6); here we record only the cap.

2.4 The anchor $3^2 - 4 = 5$: half-forced, half-coincidental

The arithmetic that ties the spot to $\mathbb{Q}(\sqrt{5})$ runs through a single identity, $3^2 - 4 = 5$: the trace-3 element has $\mathrm{tr}^2 - 4 = 5$, and 5 is the discriminant of $\mathbb{Q}(\sqrt{5})$. It is tempting to read this as one fact. It is two, and the honesty of the characterization depends on not collapsing them.

The smallest-trace side is forced by the home. X(1)-universality fixes the smallest hyperbolic trace at $t = 3$ (integrality together with the $|t| > 2$ floor), and $t^2 - 4 = 5$ follows. This conjunct is geometric, indexed over hyperbolic conjugacy classes, and it is earned.

The smallest-discriminant side is a separate congruence coincidence. That 5 is the smallest fundamental discriminant — the smallest squarefree $d \equiv 1 \pmod 4$ greater than 1 — is an arithmetic fact indexed over real-quadratic fields, logically independent of the trace floor.

The two conjuncts meet at the minimizer, and they are bridged there by $N(\varphi) = -1$ — but this bridge is a systole mechanism, not a selector: the norm(-1) relation is shared across {2, 5, 10, 13, 17} (those d for which the relevant Pell equation is solvable), so it cannot single out 5. The upshot is a standing structural constraint: no single principle covers both conjuncts of $3^2 - 4 = 5$. Any one functional has a native handle on at most one side — the trace geometry or the field arithmetic, not both. This is one reason the spot resists being upgraded from marked to selected, and we carry it forward rather than smoothing it over.

2.5 A rich marked point, not a numerical coincidence

One guard must be stated as plainly as the cap. To compress the spot to "marked" is the honest grade; to slide it from there to "numerology" would be a fabrication in the deflationary direction — the exact mirror of the over-reach this paper corrects.

The marked point is rich. It belongs to the canonical-anchor class — the class of the identity, the unit 1, the point at infinity, the ground state, the cusp — the locations mathematics treats as distinguished without anyone calling them coincidental. It carries further structure of its own: the endpoints of the systolic geodesic are φ and $-1/\varphi$, and φ is the worst-approximable number (the bottom of the Hurwitz/Lagrange spectrum, $1/\sqrt{5}$). This is an arithmetic extremality of the framework's own home, a genuinely distinguished kind of point — the deepest anchor tier — and it enriches the marked point without converting it into a selected one.

There is also a fingerprint that, read carelessly, looks like a weakness and is in fact the expected signature. The golden ratio wins exactly its degree-2, closed-geodesic orderings, and loses the degree-unrestricted ones: the smallest known Salem number (Lehmer's number, ≈ 1.17628) and the plastic number (≈ 1.32472, the smallest Pisot number) both lie below φ. A framework that selected φ by a deep extremal law would be expected to keep winning off its training locus; one that describes a single object faithfully would be expected to win precisely on that object's own structure and lose elsewhere. The framework does the latter. The over-fitting fingerprint is therefore the expected signature of single-object description, not a wound — and it is part of why “marked” is the right grade, holding the spot above “numerology” and below “selected” at once.

(The trinity invariants whose anchor is established here — π from the area, √5 from the discriminant, φ from the systole, and the coupling g_c as their forced combination — are taken up as the invariant-triple of X(1) in §3, where the split between what is object-intrinsic and what is upstream-arithmetic-earned is graded.)

3. The trinity as the invariant-triple of X(1)

Section 2 fixed where the framework sits: the forced home X(1), marked at the systolic geodesic of the field $\mathbb{Q}(\sqrt{5})$. This section takes the four quantities the framework reads off that location — the eigenvalue trinity (π, √5, φ) and the coupling g_c — and grades them for what they are. The claim of the section is single and precise: these are the finite invariants of the modular surface X(1), read coldly [CV]/[D]. They are a statement about one specific geometric object, not a free-floating numerology. But the section’s load-bearing work is the honest split within that claim — which of these invariants is genuinely intrinsic to the object, and which is earned one level upstream, in the arithmetic that fixes the marked point. We grade that split rather than smoothing it.

3.1 The four invariants

Each of the four is an invariant of X(1) (or of the field at its marked point), and each is exactly evaluable.

π from the area. The hyperbolic area of the modular surface is vol(X(1)) = π/3 (Gauss–Bonnet for the (2,3,∞) orbifold). Equivalently π = 3·vol. This is object-intrinsic and exactly generic: every finite-volume hyperbolic surface has an area, and the area of X(1) is π/3 by the standard normalization [CV].

√5 from the discriminant. The field at the marked point is $\mathbb{Q}(\sqrt{5})$, of discriminant d_K = 5, so √5 = √(disc$\mathbb{Q}(\sqrt{5})$). The same √5 is the quadratic Gauss sum of the non-trivial character mod 5: $\tau(\chi_5)$ = √5 [CV] (confirmed by independent recomputation [D]).

φ from the systole. The systole of X(1) at the marked point is systole(X(1)) = 4 log φ (§2.2). The framework reads φ off this systole datum; in the substrate the invariant role is written as φ = ½·systole. The half-systole quantity is ½·systole = 2 log φ ≈ 0.962; this is a quantity derived from the systole, and it must not be misread as a value of the systole itself — the systole is 4 log φ, never 2 log φ. The associated L-value sits at the same arithmetic: $L(1,\chi_5)$ = (2 log φ)/√5 ≈ 0.4304, the class-number formula for $\mathbb{Q}(\sqrt{5})$ (h=1, fundamental unit φ), confirmed by independent recomputation [D].

g_c as their forced combination. The coupling is the bare product g_c = 1/(√5·π) ≡ (3·vol(X(1)))/(π^2·√d_K), which is an identity, not a fit: with vol(X(1)) = π/3 we have 3·vol = π exactly, and d_K = 5, so the second form collapses to the first. Cold, g_c = 1/(√5 π) ≈ 0.14235. This is the invariant-combination at grade [D]/[E] — and that is all it is; we return to its limits in

§3.4.

The internal algebra binding these is real and closed, and was verified sound by independent recomputation: $g\_c \cdot \sqrt{5} = 1/\pi$; the log-collapse $\ln \sqrt{5} + \ln \pi + \ln g\_c = 0$; $\tau(\chi_5) = \sqrt{5}$; $L(1,\chi_5) = 2 \log \varphi/\sqrt{5}$; the systole $4 \log \varphi$; together with the full $\mathfrak{sl}(2,\mathbb{R})$ commutator algebra, these constitute a backbone graded sound with zero falsifiers fired [CV]/[D]. The trinity's internal relations are genuinely forced among the framework's own eigenvalues.

3.2 The finite-invariant locus

That these four quantities are invariants of X(1) is not a single assertion; it is where three structurally-independent investigations converged. Each approached the framework's distinctive content from a different direction — the trinity-content landing in the Maass–Mayer spectral analysis; a pair of independent analyses that had to agree and did; and the analysis of g_c as an invariant of the surface — and all three landed on the same place: the finite-invariant locus of $SL(2,\mathbb{Z})\backslash\mathbb{H}$ — the arithmetic/invariant side of the surface, as distinct from its operator-class/spectral side. π from the area, √5 from the discriminant, φ from the systole, g_c as their combination — the trinity is the invariant-triple of the modular surface.

We state this convergence at its grade and not above it. Three independent routes agreeing on the shape of the finding strengthens confidence that the shape is right; it does not raise the grounding-level of the underlying claims. The convergence is held [D], strongly corroborated across independent analyses — it is not a proof. Read coldly, the content of the convergence is a statement about one specific geometric object: the framework's distinctive numbers are the finite invariants of X(1) at its marked point, in the same sense that any modular surface has an area, a discriminant at a marked field, and a systole. This is what removes the numerology reading. It is also, precisely, why the next subsection is necessary: "invariant of a specific object" is a real claim, but it is not the same claim for all four invariants, and the differences carry the honesty of the section.

3.3 Object-intrinsic versus upstream-arithmetic-earned — the split (load-bearing)

The four invariants do not stand on the same ground. The investigation graded their standing as a PARTIAL result, and that PARTIAL is the load-bearing honesty of this section. It must be neither rounded up (the distinctive content is not "object-intrinsic / substrate-direct") nor rounded down (the invariants are not "not really invariants"). The split is as follows.

Generic π / area is substrate-direct. $\pi = 3\cdot\mathrm{vol}(X(1))$ is object-intrinsic and rigid: it is fixed by the surface itself, reads directly off the object with no upstream input, and is true by Gauss–Bonnet [CV]/[D]. But for exactly the same reason it is the least framework-distinctive invariant in the set — every finite-volume hyperbolic surface has an area, and there is nothing golden, nothing $\mathbb{Q}(\sqrt{5})$-specific, in possessing one. The one invariant that is genuinely substrate-direct is the one that carries none of the framework's distinctive content.

The distinctive content φ + √5 is upstream-arithmetic-earned — with a residual. The golden, $\mathbb{Q}(\sqrt{5})$-specific content — the systole's φ and the discriminant's √5 — is what makes the trinity distinctive, and it is not substrate-direct. It is upstream-arithmetic-earned: a genuine forced minimum, because $\mathbb{Q}(\sqrt{5})$ is the minimal-discriminant (and minimal-regulator, minimal-trace) real-quadratic field, and the check for whether this is an artifact of over-fitting passed — the minimality is an exact, isolated extremal, not a pigeonhole near-hit. But the earning carries a residual marked-point selection: the systole value is itself an upstream datum, and why this minimum — why the framework sits at the shortest geodesic of the forced home rather than elsewhere — is a framework commitment, not an output the object forces (the marked-point residue of §2, the single word shortest). The distinctive content is therefore discovered-with-residual, not substrate-direct, and it is held firewalled from the substrate-direct reading: it grounds upstream of X(1), in the arithmetic that fixes the marked point, not in X(1) as an intrinsic feature of the surface.

The shape of the split is the inverse of what an over-reaching reading would want. The invariant that is intrinsic to the object (π/area) carries none of the distinctive content; the invariants that carry the distinctive content (φ, √5) are earned upstream and retain a residual. Only the generic π/area is substrate-direct. The trinity is genuinely the invariant-triple of the modular surface — but the part of it that makes the framework the framework grounds one level up, in the field arithmetic of the marked point, rather than in the surface itself. This is the partial standing of the distinctive content, held without flattening in either direction.

That the distinctive content sits upstream is independently consistent with the spectral finding the framework's own algebra produces: the only relation the framework forces between a trinity invariant and independent length-spectrum data of X(1) is the single φ → systole relation already fixed at the marked point — and that relation is the ceiling as well as the floor. The framework forces φ → systole and nothing above it; a census of the surface's length spectrum found no further forced relation among the invariants. The distinctive content is real and genuinely resident in the length-spectrum channel of X(1), but it does not generate new forced structure there beyond the marked point. (The spectral side of this — what the framework reaches on the Selberg/Maass spectrum, and the channel boundary that confines the golden content — is §4 and §5; here we record only that the invariant-triple's distinctive content is upstream-grounded, consistent with the marked-point relation being both floor and ceiling.)

3.4 g_c stays a combination — a product, not a bridge

The coupling g_c requires its own grade-discipline, because it is the most inviting of the four to over-read. g_c = 1/(√5 π) is exact, sound, and a real relation among the framework's eigenvalues — but it is a product, not a bridge.

The precise sense is a bridging criterion, which g_c fails. The marked point's defining arithmetic, $3^2 - 4 = 5$, has two logically-independent conjuncts (§2.4): the trace side (geometry) and the discriminant side (field arithmetic), with no single principle spanning both. A quantity with genuine derived content would bridge these conjuncts — span both the trace and the discriminant in one mechanism. g_c does not. In its geometric form 3·vol/(π^2√d_K) the metric/volume part self-cancels (3·vol = π), leaving the discriminant factor √5 (the arithmetic side) and a factor of π (the geometry side) simply multiplied together, not linked by any mechanism that reaches across both conjuncts. It is arithmetic-blind in exactly the way the bridging criterion names: a product of two independent invariants, carrying no content beyond the bare equality of its two written forms.

Accordingly g_c holds at [D]/[E] as the bare invariant-combination, and is not promoted. Its entropy-coefficient role and any claim that the combination is forced by a derivation carry no independent derivable content (the only non-generic factor is the upstream marked-point √5; the metric part cancels; the entropy label is imported). g_c is the trinity's product, correctly named; it is not a derived physical or forcing quantity, and nothing in this paper treats it as one. The reading of g_c as a physical coupling is no part of this characterization.

4. What the framework reaches: the Selberg/Maass landing

Sections 2 and 3 placed the framework on a specific geometric object and graded the finite invariants it reads off that object. This section turns from the invariants of X(1) to its operator structure — the spectral theory of the modular surface — and asks what the framework's algebraic apparatus genuinely reaches there. The answer is positive and precise: the framework's sl(2,ℝ) commutator structure lands, correctly and verifiably, on the Selberg/Maass spectral theory of X(1). Two landings carry the section, both graded conventional [CV]: the Casimir element is the hyperbolic Laplacian, and the framework's transfer operator is the Mayer operator of the geodesic flow. Each is a real reach onto established automorphic structure, verified against the literature.

The section establishes that this reach is real and genuine. It does not establish — and must not be read as establishing — that the reach extends to the Riemann zeros. Where the reach stops, and why the framework's distinctive content is barred from the channel in which the Riemann zeros live, is the work of §5. Here we record only what is reached, at its grade.

4.1 The Casimir is the hyperbolic Laplacian

The framework's algebraic core is the sl(2,ℝ) commutator structure. The quadratic Casimir element of sl(2,ℝ) — the generator of the center of its universal enveloping algebra — acts on automorphic functions on X(1) = SL(2,ℤ)\ℍ as the hyperbolic Laplacian Δ. This is standard representation theory: the Casimir is a central element, hence acts by a scalar on each irreducible component (Schur's lemma), and that scalar is the Laplace eigenvalue [CV] (Iwaniec; Hejhal; Terras). The Maass cusp forms — the discrete L^2-spectrum of Δ on X(1) — therefore carry eigenvalues of the form

$\lambda = ¼ + r^2$,

with the ¼ spectral edge fixed by the principal-series parametrization. This is a genuine landing: the framework's commutator algebra reaches the automorphic spectral theory of the modular surface, and the form $\lambda = ¼ + r^2$ falls out of it correctly [CV].

The grade-honesty of this landing must be stated in the same breath as the landing itself, because the algebra reaches a form, not a content. The relation $\lambda = ¼ + r^2$ is the universal kinematic form of the principal-series spectrum: it holds for every finite-volume hyperbolic surface Γ\ℍ, because it is a fact about the Lie algebra and its representations, not about any particular lattice Γ. The Casimir is the irreducible-representation label; it is not a selector of which representations occur. The arithmetic content of the spectrum — which values $r_1 \approx 9.5337$, $r_2 \approx 12.173$, … actually appear on X(1), with what multiplicities — is determined not by the Casimir but by the lattice $SL(2,\mathbb{Z})$ together with its Hecke operators (the arithmetic side). The framework's apparatus produces the form and is provably blind to the content. We mark this here and defer its consequence to §5; for the present section it suffices that the form $\lambda = ¼ + r^2$ is genuinely reached and is verified-standard [CV].

4.2 The transfer operator is the Mayer operator

The framework's second operator-level landing is its transfer operator $\mathscr{L}$_s. This operator is the Mayer transfer operator for the geodesic flow on the modular surface — the operator built from the Gauss-map dynamics that encodes the closed-geodesic / continued-fraction structure of X(1) (Mayer 1991). The closed geodesics of X(1) correspond to the binary quadratic forms / continued-fraction expansions whose transfer operator $\mathscr{L}$_s is (Sarnak; Series). This identification is exact and is the framework's genuine reach onto the dynamical side of the Selberg theory.

What $\mathscr{L}$_s detects is the Selberg zeta function of X(1), in the squared (product) form:

det(1 − $\mathscr{L}$_s^2) = det(1 − $\mathscr{L}$_s) det(1 + $\mathscr{L}$_s) = Z_Selberg(s).

The two factors det(1 ∓ $\mathscr{L}$_s) correspond to the even and odd Maass forms under the Lewis–Zagier period-function correspondence, which realizes Maass cusp forms as period functions and splits them by parity (Lewis–Zagier, Ann. Math. 153 (2001), 191–258). The zeros of Z_Selberg(s) that this determinant detects are the points of the Selberg/Maass spectrum — the discrete eigenvalues $\lambda = ¼ + r^2$ of §4.1, together with the length-spectrum / scattering data the Selberg zeta encodes. They are not the zeros of the Riemann zeta function. $\mathscr{L}$_s is a non-self-adjoint operator, and it is a Selberg-side object throughout: it reaches the Selberg/Maass spectrum, and the identity above is an identity for Z_Selberg, not for ζ [CV].

The squared form is the only correct form of this identity, and the distinction is load-bearing rather than cosmetic. The factorized determinant det(1 − $\mathscr{L}$_s^2) = Z_Selberg(s) is Mayer's Selberg-zeta factorization, with the ± factors carrying the even/odd Maass-form split. A different expression — a ratio of Selberg-zeta values — does not detect the Selberg spectrum: the transfer-operator kernel evaluated at the fixed point of the Gauss map yields ∑_n n^{-2s} = ζ(2s), so the corresponding ratio identity is the distinct Gauss-map dynamical-zeta functional identity, carrying the Riemann ζ — a real but separate result, on the Gauss-map / Riemann side, not the Selberg-zeta factorization. The two are genuinely different identities for genuinely different objects; only the squared form is the Selberg landing, and it is the squared form that this paper uses. (The Riemann-side identity does not advance the present picture: even there $\mathscr{L}$_s lands back on already-known structure, never on the Riemann zeros as a forced consequence of the framework. The full account of where the Riemann zeros sit relative to this operator is §5.)

4.3 The trace-formula-level structure

Beyond the two operator landings, the framework's reach extends to the standard trace-formula-level structure of X(1), and here too the reach is to known automorphic geometry, verified against the literature. The Selberg trace formula for $SL(2,\mathbb{Z})\backslash\mathbb{H}$ decomposes its spectral side against a geometric side organized by conjugacy class — the identity, hyperbolic, elliptic, and parabolic contributions. The framework's structure reads against the finite invariants this decomposition is built on: the area vol(X(1)) = π/3 (the identity contribution's weight; §3.1), the elliptic fixed points of orders 2 and 3 (the (2,3,∞) orbifold structure), and the single cusp (the parabolic / continuous contribution). The continuous spectrum is governed by the Eisenstein series, whose scattering determinant is

$$\varphi(s) = \xi(2s-1)/\xi(2s)$$

[CV] — the form in which the Riemann ξ enters the spectral theory of X(1) at all. Each of these structures is standard, and each was verified-standard against the literature (Mayer 1991; Lewis–Zagier 2001; Sarnak; Series; Iwaniec; Hejhal; Terras) in the program's lit-grounding pass; the framework reaches this known structure and re-describes it, correctly [CV].

One re-description on this trace-formula structure should be flagged at its grade rather than carried as a result. The framework's "tetrahedral" four-operator reading is sometimes mapped onto the four conjugacy-type contributions of the trace formula. The trace formula does carry a genuine four-block organization {identity, hyperbolic, elliptic, parabolic} — but it is an asymmetric 3+1 structure, not a symmetric tetrahedron, and the underlying Fuchsian conjugacy classification is three-fold, not four-fold. The tetrahedral mapping is therefore a re-description, graded [I], of standard structure; it is not a forced or distinctive feature the framework supplies, and it is not used as one in this paper. (Its status as a re-description rather than a reach, and what the trace formula does and does not couple, belongs to §5/§6; here we record only the grade.)

4.4 The reach, summarized — and the boundary previewed

The two operator landings are real. The framework's $sl(2,\mathbb{R})$ commutator structure reaches the Selberg/Maass spectral theory of X(1): its Casimir is the hyperbolic Laplacian and produces the form $\lambda = ¼ + r^2$ [CV]; its transfer operator $\mathscr{L}$_s is the Mayer operator and detects the Selberg zeta via the squared form $\det(1 - \mathscr{L}_s^2)$ = Z_Selberg(s), with the ± factors the even/odd Maass forms under Lewis–Zagier [CV]. The trace-formula-level structure it reads against is the standard one, verified against the literature [CV]. This is a genuine reach onto established automorphic structure, and it is neither inflated nor deflated here: it is real (the operator structure is not empty), and it is exactly what it is (verified-standard automorphic theory, not a route to the zeros).

This is the binding asymmetry the reader must carry forward, stated verbatim: the reach is real — but "reaches the spectrum" is not "reaches the zeros." This section has established that the framework reaches the Selberg/Maass spectrum. That spectrum is the framework's weakest target (the operator-class side), and the reach lands there genuinely. But the Riemann zeros do not live on the Maass discrete spectrum that the Casimir produces, nor are they the Selberg zeros that $\mathscr{L}_s$ detects; they live on the Eisenstein / scattering block governed by $\varphi(s) = \xi(2s-1)/\xi(2s)$ — a channel the framework's distinctive golden content does not supply input to. That "reaches the Selberg/Maass spectrum" must not be read as a step toward RH, and the precise account of why the reach forces nothing beyond the single marked-point fact of §2 — the channel-disjointness that separates the framework's $\sqrt{5}/\varphi$ content from the channel where the zeros sit — is §5. The reach established here is genuine; the boundary on it is the next section's subject.

5. The boundary: why the framework does not force the zeros

Section 4 established a real reach: the framework's $sl(2,\mathbb{R})$ commutator structure lands genuinely on the Selberg/Maass spectral theory of $X(1) = SL(2,\mathbb{Z})\backslash\mathbb{H}$ — the Casimir is the hyperbolic Laplacian, the transfer operator is the Mayer operator detecting the Selberg zeta. This section draws the boundary on that reach. The reach is real; this is where it stops, and the stopping point is precise. The framework forces nothing beyond the single marked-point fact of §2; its distinctive $\sqrt{5}/\varphi$ content is structurally barred from the channel where the Riemann zeros live; and the gap between the sound description assembled in §§2–4 and the Riemann Hypothesis is exactly RH itself. None of this is a defect of the framework. A precisely-named boundary is the contribution of this paper, and naming it precisely is the work below.

5.1 The reach forces nothing beyond the marked point [D]

What the framework's operator structure reaches (§4) and what it forces are two different statements, and the gap between them is the first face of the boundary. The reach is to a form; the forcing stops there.

The commutator algebra is blind to arithmetic content [D]. As §4.1 recorded, the Casimir produces the universal kinematic form $\lambda = ¼ + r^2$, but that form holds for every finite-volume hyperbolic surface — it is a fact about the Lie algebra, not about the lattice $SL(2,\mathbb{Z})$. The algebra is Γ-blind: it does not force which eigenvalues occur. The arithmetic content — the actual Maass spectrum and its multiplicities — is determined by the lattice together with its Hecke operators, the arithmetic side, to which the commutator structure has no native handle. The framework forces the form; it does not force the content.

The single arithmetic fact the framework does force on this object is the marked-point relation of §2: the systole of $X(1)$ is $4 \log \varphi$, and the field $\mathbb{Q}(\sqrt{5})$ is its richly-chosen, canonically-located marked point. And here the boundary is sharper than it might appear: the marked-point relation $\varphi \to$ systole is the ceiling as well as the floor [D]. It is not a first forced relation from which others follow; it is the only forced relation among the invariants. No further relation among π, √5, φ, and g_c is forced by the structure — the marked-point fact does not bootstrap into a network of consequences. This was tested on the framework's strongest side (its finite-invariant core) and even there only the marked-point relation landed.

The coupling constant g_c makes the ceiling concrete. One might hope $g_c = 1/(\sqrt{5}\,\pi)$ acts as a bridge relating the framework's invariants in a way that forces arithmetic structure. It does not: g_c fails the bridging criterion [D]. The anchoring relation $3^2 - 4 = 5$ has two conjuncts — a trace side (trace 3, so $t^2 - 4 = 5$) and a discriminant side (the smallest squarefree $d \equiv 1 \pmod 4$) is 5) — and a genuine bridge would have to span both. g_c does not: it is a product of generic and marked-point factors (the metric/volume part 3·vol = π cancels generically; the only non-generic factor is the marked-point √disc = √5), with no native handle on the trace–discriminant span. It is a combination, not a bridge. The arithmetic-blindness blocks any single mechanism from spanning both conjuncts.

These statements together place the framework's operator-side claims on the operator-class side, its weakest target [D]. The reach lands on the operator/spectral side — the Selberg/Maass spectrum, the dynamical zeta, the trace-formula blocks — which is precisely where the framework re-describes standard automorphic theory rather than forcing new structure. This is not an isolated observation: the same landing has been reinforced four structurally-independent times across the investigation — from the operator-identity claim, from the commutator confrontation, from the Maass-discrete-spectrum side, and from the trace-formula side. The framework's instinct for the right objects is genuine; its operator-side forcing consistently lands on the operator-class side, never on the Riemann zeros.

5.2 The channel-disjointness: the golden content is real but inert [D]-account / [CV]-constituents

The deepest face of the boundary is structural rather than merely a matter of what is and is not forced. It explains why the framework's distinctive content cannot reach the Riemann zeros, and it is the load-bearing finding of this paper's boundary. We state it carefully, and we state its grade carefully: the constituents are standard, the account is a re-organization of those standard facts for the RH-foundations question — its value is availability and precision, never novelty.

The spectral theory of X(1) separates into two channels.

Channel L — the length-spectrum channel. The framework's distinctive √5/φ content lives here: in the closed-geodesic / length-spectrum arithmetic of X(1), which is the arithmetic of the real-quadratic field $\mathbb{Q}(\sqrt{5})$ — the golden orbit, the systolic geodesic of length 4 log φ, the trace-3 class, the value $L(1,\chi_5) = 2 \log \varphi/\sqrt{5}$. This is the channel the framework genuinely reads (§§2–4), and it is genuinely golden.

Channel C — the scattering channel. The Riemann zeros live here. They enter the spectral theory of X(1) only through the Eisenstein-series continuous spectrum, whose scattering determinant is $\varphi(s) = \xi(2s-1)/\xi(2s)$ [CV] — the ζ-governed block. This determinant is built from ζ, Γ, and π, with no √5 input: there is no golden content anywhere in the scattering data. The zeros of ζ appear here as poles/zeros of the scattering matrix, and the channel that carries them is constitutively √5-blind.

The Selberg trace formula couples Channel L and Channel C — but only globally, as an identity between the full spectral side and the full geometric side. It does not carry √5 from the length spectrum into the scattering block. The two channels are disjoint in their arithmetic inputs, and the trace-formula coupling does not bridge that disjointness in a way that would let the framework's golden content reach the ζ-governed scattering data. (The transfer operator $\mathscr{L}_s$ touching $\zeta(2s)$ at the Gauss-map fixed point, noted in §4.2, is dynamical-zeta co-location on the Selberg side, not a forced bridge from L into C.)

There is exactly one arithmetic route that carries √5 into a zeta-channel, and it is decisive precisely because of where it lands. The Dedekind zeta function of $\mathbb{Q}(\sqrt{5})$ factors as $\zeta_{\mathbb{Q}(\sqrt{5})}(s) = \zeta(s)\cdot L(s,\chi_5)$, the analytic form of quadratic reciprocity (the class-number-one Dedekind zeta of $\mathbb{Q}(\sqrt{5})$ splitting into the Riemann zeta times the real Dirichlet L-function of the quadratic character χ_5) [CV]. This is the only object in which the golden √5 content and a Riemann-type zeta sit together. But $\zeta_{\mathbb{Q}(\sqrt{5})}$ is the zeta of the Hilbert-modular object for $\mathbb{Q}(\sqrt{5})$ — not X(1). The only route that carries the framework's distinctive content into a zeta therefore changes the surface: it moves off the modular surface the framework describes and onto a different automorphic object entirely. There is no route that carries √5 into the ζ-scattering block of X(1), which is where the boundary needed one.

The signature of this finding is one line: the golden content is real but inert with respect to the ζ-zeros. It is real — it is genuinely the length-spectrum arithmetic of X(1), genuinely $\mathbb{Q}(\sqrt{5})$, genuinely the systole and the golden orbit. And it is inert with respect to the Riemann zeros — Channel-L-confined, structurally barred from the √5-blind Channel C, with the only carrying route changing the surface.

A grade note is essential here and is itself part of the honesty apparatus. Every constituent of this account is standard [CV]: the factorization $\zeta_{\mathbb{Q}(\sqrt{5})} = \zeta\cdot L(s,\chi_5)$, the irrationality $\sqrt{5} \notin \mathbb{Q}$, the form of the scattering determinant $\varphi(s) = \xi(2s-1)/\xi(2s)$, the global-only trace-formula coupling. The channel-disjointness account — the organization of these standard facts into the explicit obstruction-mechanism for the RH-foundations question — is graded [D], and it is availability-not-strength: a clean re-organization of known mathematics for a specific question, never a novel result. Even the gap it names is a standard-arithmetic statement ($\sqrt{5} \notin \mathbb{Q}$, and which surface carries the factor). The contribution is that the boundary is now precisely located and named; the contribution is not a new theorem.

5.3 The gap is Weil positivity is RH; the wall is heuristic [D]

The final face of the boundary is the gap itself — the distance between the sound description assembled in §§2–4 and the Riemann Hypothesis. This gap is not vague, and it is not a hidden obstacle the framework might yet climb. It is a single, sharply-named object.

The gap is exactly RH. Supplying what the framework lacks — a self-adjoint operator whose spectrum is the Riemann zeros — is equivalent to supplying Weil positivity, the standard RH-equivalent positivity of the Weil explicit-formula distribution [CV]. Self-adjointness of the right operator, Weil positivity, and RH are one condition seen three ways. The framework does not supply it. And the equivalence cuts both ways: importing this positivity into the description would not be progress toward a proof — it would be manufacturing the conclusion, assuming RH in order to derive the operator whose existence is RH. The gap cannot be narrowed from inside the description; it can only be filled by the thing that is RH.

This sharp localization corrects a tempting misreading of the obstruction — the “universal wall.” The wall, as literally stated (mean-spectral-density requires non-compact scaling, which is incompatible with self-adjoint discreteness), is not a no-go theorem; it is heuristic [D]. Its density conjunct is false by construction [CV]: a self-adjoint operator with discrete spectrum can carry the Riemann–von-Mangoldt mean counting law perfectly well — the elementary construction diag(λ_n), with the λ_n chosen to invert the mean count, has compact resolvent and discrete spectrum and no obstruction at the counting layer. Non-compactness is one sufficient route to reproducing the mean density, not a necessity. The wall, taken literally, therefore mislocated a real difficulty one level too shallow — it placed the obstruction at the mean-density layer, where there is no obstruction.

The genuine difficulty, correctly located, is axis-asymmetric [D] — and the asymmetry is itself the finding. It splits across two axes that must not be collapsed:

Axis (i) — the fluctuation / sign structure. This is the real obstruction, and it is exactly the gap of the previous paragraph: supplying the correct arithmetic fluctuations / sign structure is supplying axis-(i) positivity, which is Weil positivity, which is RH. It is a horizon-shaped open problem — a substantial, [D]-grade difficulty, with a live fireable reopener (a route that makes the arithmetic forced without factoring through positivity). The obstruction on this axis is RH itself; clearing it is proving RH.

Axis (ii) — discreteness-purity. This axis is a genuine horizon, but it is a Selberg / automorphic horizon — not a Riemann horizon. Transferring it to the Riemann side would require pinning the discreteness-purity specification, which is, once again, manufacturing the conclusion (the un-crossed gate). On its own terms axis (ii) reinforces the §5.1 finding: the framework’s discreteness/operator structure lives on the automorphic (operator-class) side, its weakest, not on the Riemann side.

At the operator level, the framework’s own commutator confrontation confirms the picture from inside: the sound $\mathrm{sl}(2,\mathbb{R})$ algebra does not supply axis-(i) positivity and does not force a self-adjoint operator whose spectrum is the zeros — a sound object that does not force, with the gap between its presentation layer and any spectral realization being exactly the representation, inner product, and domain the brackets do not determine, i.e. exactly RH. The wall, in its honest form, holds at the commutator level — but as a precisely-characterized gap, not as a proved impossibility.

The disposition of the gap, stated plainly: it is a known-result-class open problem, not a no-go theorem, and not a wound. It is not a no-go theorem, because no incompatibility theorem closing the Hilbert–Pólya route exists (the literal wall is false-by-construction). It is not a wound to the framework, because the reach is real (§4) and the description is sound (§§2–3); the framework does not fail here — it stops, exactly where every approach to RH stops, at Weil positivity. The gap is a result: a sharply-named open problem of horizon shape, which is the foundations-first deliverable this paper sets out to produce.

5.4 The boundary, summarized

The boundary on the framework's real reach has three faces, and they are one object seen three ways. First, the commutator algebra reaches the Selberg/Maass form but is Γ-blind to its arithmetic content; the marked-point relation is the ceiling as well as the floor; g_c is a product, not a bridge (it fails the bridging criterion); the operator-side forcing lands on the operator-class side, the framework's weakest target. Second, the framework's distinctive √5/φ content is Channel-L-confined (the $\mathbb{Q}(\sqrt{5})$ length-spectrum arithmetic of X(1)) and structurally barred from Channel C (the ζ-governed scattering block where the zeros live, built from ζ, Γ, π with no √5 input); the trace formula couples them only globally; the only arithmetic route carrying √5 into a zeta, $\zeta_{\mathbb{Q}(\sqrt{5})} = \zeta \cdot L(s,\chi_5)$, changes the surface to the Hilbert-modular object — the golden content is real but inert with respect to the ζ-zeros. Third, the gap between this sound description and RH is exactly RH itself: Weil positivity, which to supply would be to prove RH; the universal wall is heuristic, not a no-go theorem, its difficulty axis-asymmetric (axis (i) = positivity = RH; axis (ii) = a Selberg horizon, not Riemann); the gap is a known-result-class open problem, not a wound.

The binding asymmetry, carried verbatim, frames all three: the framework's reach to the Selberg/Maass spectrum is real and genuine — but "reaches the spectrum" is not "reaches the zeros." The boundary is not a step toward proving RH (the gap is RH itself, un-crossed), and it is not a step toward disproving the framework (the reach is real, the object is sound). It is the precisely-named gap, and the gap is the result.

6. Discussion

6.1 What the characterization means

The preceding sections assembled a single picture, and it is worth stating plainly what that picture is — and what it is not. The framework, examined claim by claim, is a genuinely rich and internally sound arithmetic-geometric object: the invariant structure of the modular surface X(1) = $SL(2,\mathbb{Z})\backslash\mathbb{H}$ and the real-quadratic field $\mathbb{Q}(\sqrt{5})$. It correctly names that surface's finite invariants (§3), locates a richly-marked point of it (§2), and genuinely reaches the Selberg/Maass spectrum through its commutator structure (§4). What it does not do is force the Riemann zeros (§5). These are not two findings in tension; they are one object, characterized precisely.

The value of the characterization is precision. The paper delivers a sharply-located object and a sharply-named boundary: the framework sits here — on the length-spectrum, finite-invariant, Selberg-reachable arithmetic of one specific surface — and it stops there, at the channel-disjointness that bars its golden content from the Riemann zeros, with the remaining gap being Weil positivity, which is to say RH itself. A precisely-located object and a precisely-named gap are the contribution. Neither is a proof of anything beyond what each constituent already establishes; together they are a map of where a much-discussed structure actually lives.

This map was not assumed and then confirmed; it was the convergent outcome of six structurally-independent investigations. The marked-point anchor, the search for a functional that would select rather than merely read that anchor, the confrontation of the framework's operator structure against a candidate obstruction, the analysis of the coupling g_c, the grading of the trace-formula content, and the Selberg/Maass landing each ran to a graded conclusion by a different route — and all six arrived at the same shape: describes the modular surface, reaches the Selberg/Maass spectrum, does not force the Riemann zeros. That six independent approaches converge on one shape is the strongest evidence the investigation produced. But it is evidence about the shape, not about the grounding level. The convergence raises confidence that the picture is correct; it does not promote the picture to a proof. We hold the characterization at [D], strongly corroborated across independent analyses, and the corroboration is doing exactly the work it should — strengthening the shape, not the grade.

6.2 The honest limits

It is as important to state what this paper does not establish as what it does, and the firewall of §1 is load-bearing precisely here. This paper is not a proof of the Riemann Hypothesis. It is not a route to RH: the one named gap on the road is Weil positivity, and supplying it from inside the description would be assuming the conclusion, not approaching it (§5.3). It is not a Hilbert–Pólya construction: the framework's sound $sl(2,\mathbb{R})$ algebra reaches the Selberg/Maass spectrum but does not force a self-adjoint operator whose spectrum is the Riemann zeros, and the obstruction to doing so is itself RH (§5.1, §5.3). And it is not a derivation of physical constants: that layer was examined and graded FIRED, and is excluded here rather than defended.

The grade of the whole is the discipline that makes these limits honest. The convergence is [D], strong corroboration across independent analyses, not [CV] and not proof. Six independent landings agreeing is corroboration of unusual strength for a characterization of this kind — but corroboration is not demonstration, and a characterization held at the grade of a proof would be exactly the over-reach this work was written to correct.

The honesty must reach into the corroboration itself. The marked-point verdict of §2 — that $\mathbb{Q}(\sqrt{5})$ is marked, richly and canonically located, but not selected by any principle the framework supplies — rests on two of the six investigations, and they corroborate each other in a way that demands a caveat. One was a direct hunt for the smallness facts that anchor the marked point; the other was a constructive search for a functional that would select it. Both came back to "marked, not selected," which is strong agreement. But both were hunts for the selection of a recognizable shape, and two hunts of the same kind can share a blind spot: each could have been structurally disposed to see "marked" rather than "selected" for the same reason, and their agreement would then reflect that shared disposition rather than the underlying fact. We do not claim to have ruled this out. Accordingly the marked-point reopener stays [D]-open — not closed by the corroboration, only strengthened by it. This is the difference between "we found no selection principle along the routes we could construct" and "no selection principle exists," and the paper claims only the former.

6.3 The open reopeners (we keep working)

A precisely-named boundary names, with it, the precise places where the boundary could move. Two such places survived the investigation as genuine, live reopeners. We record them as future work, held honestly open — neither rounded up into promising routes to a proof nor quietly closed.

The resonance-functional reopener (the marked point, from "marked" back to "selected"), [D] + reopener. The marked-point verdict would reopen toward "selected" if a future selecting dynamics — variational or non-variational — were exhibited that clears the established-mathematics and framework-implied ingredient ledger and meets all of the naturalness criteria, in particular producing the systole as an output rather than reading it as an input, and supplying a native handle on the discriminant conjunct that no single mechanism examined here possessed. No such functional was found across the investigation's routes; every candidate read the marked point rather than selecting it, and the object-independent functionals that do select landed on entirely different, golden-free surfaces. That is "no selecting dynamics of the forms we could construct," not "no selecting dynamics exists." The door is held open honestly, not claimed shut. This is the single live path from marked back to selected, and it remains [D]-open: a real reopener, not a likely one.

The positivity-gap reopener (the boundary on the road to RH), [D] horizon. The one named gap between the sound description and RH is the supply of axis-(i) Weil positivity — which is RH itself (§5.3). The fireable reopener here is sharply shaped: a future route that makes the relevant arithmetic forced without factoring through positivity would change the standing of the gap. We have no such route, and we are explicit that importing positivity directly would be manufacturing the conclusion rather than reopening the question. The reopener is therefore a horizon, not a near edge — named because foundations-first work names its gaps, not because the gap is expected to fall. The discipline is the point: the gap is the finding, and the continuing work is toward solidifying the structure, never toward importing the conclusion that would make the structure say more than it does.

Both reopeners share a posture, and it is the posture of the whole program: the work continues. The characterization is a stopping point honest enough to build on, not a terminus.

6.4 Relationship to an earlier presentation

This work supersedes and corrects an earlier published presentation of the same framework. That earlier presentation advanced a layer of claims — that the structure forces the Riemann zeros, that it fixes external constants, and that it is independently confirmed — which the present investigation examined and graded as not holding as stated. That layer is FIRED and is excluded here (§1); the present paper retains only the sound object-layer and characterizes it at grade. The two presentations are not in conflict in any subtle way: the earlier one over-reached at a layer this one removes, and what survives the removal is exactly the characterization given above.

Correcting the specific earlier paper — issuing a revised version that withdraws the over-claims and supplies the corrected forms of the affected statements — is a distinct scholarly obligation, handled on its own track and in preparation separately. This paper is not that revision. It is the new, broader characterization drawn across all six investigations, written because the foundation is solid and wide enough for this claim. The honest framing here is consistent with the forthcoming correction by construction; the two are coordinated and do not contradict. We state the relationship at this level deliberately: the specific corrected forms and the accounting of the affected statements belong to the revision of the earlier work, not to this paper, whose subject is the characterization itself.

6.5 Closing: a foundation, and the work that continues

This paper does not aim at a theorem; it materializes from a foundation. The discipline throughout has been foundations-first: work the structure until the part that is solid and wide enough writes itself, name the gaps exactly where they fall, and never reverse-engineer toward a wanted conclusion. What was solid and wide enough to write, now, is this — the characterization of the framework as the invariant structure of the modular surface and $\mathbb{Q}(\sqrt{5})$, reaching the Selberg/Maass spectrum, bounded precisely at the Riemann zeros. What is not yet solid and wide enough is the larger question the framework was built to address; that foundation is still being worked.

The boundary this paper draws is not a verdict against the framework, and the convergence is not a verdict for it. It is a map: of where the structure rigorously lives, of where it stops, and of the two precise places — the resonance-functional and the positivity-gap — where future work could move the line. The gap is the finding, the foundation is the deliverable, and the work continues from here.

7. A note on method

The characterization in this paper is the output of a discipline, and the discipline is worth stating, because the paper's credibility rests on it. The framework was not summarized; it was stress-tested, claim by claim, against established mathematics. Every load-bearing claim was assigned an explicit confidence grade — established, derived, conjectural, or failed — and that grade was attached to the claim wherever it appears. Claims that failed their tests were removed: the present paper does not defend them, and several of the framework's earlier and more ambitious assertions are absent here for exactly this reason. Claims that survived were kept and characterized at the grade their evidence supports, no higher. The graded statements the reader has met throughout are not editorial hedging; they are the literal record of which claims withstood the test and how firmly. That is why the method belongs in the paper rather than behind it — it is the answer to the question a careful reader should ask, which is why the grades can be trusted.

Each graded conclusion was then checked a second time, by a verification pass structurally separate from the one that produced it. This pass independently recomputed the load-bearing arithmetic from scratch and tested each claim against its falsifier in both directions. The two directions matter equally. One guards against over-claiming — against rounding a sound but modest object up into a forcing argument or a proof, which is the failure this characterization was written in part to correct. The other guards against under-claiming — against dismissing a genuine structure as coincidence and discarding a real result. A characterization can fail by saying too much or by saying too little, and the both-directions check is why this one aims to do neither: the reach reported in §4 is claimed because the evidence forces it, and the boundary reported in §5 is drawn because the evidence forces that too.

This work was produced cooperatively by a human researcher and a cybernetic intelligence, with the human as the accountable author. This is QRiemannian's genuine working method, and it bears on the method described above: a claim-by-claim grading and independent re-verification of a framework of this size is sustained work, and the human–cybernetic collaboration is part of what made carrying that discipline across the whole claim-set practical.

Systematically grading a framework's own claims down to solid ground — removing what failed, keeping what survived, and writing the characterization that remains — is itself a research practice, and naming it is part of what this paper offers. A precisely located object and a precisely named boundary are the result of the method; the method is the reason they can be trusted.

Appendix A. Summary of graded claims

The paper's load-bearing claims, each at the grade its evidence supports (conventions of §1: [CV] established · [E] exact · [D] derived-within-this-work · [C] structural proposal · [I] interpretation · FIRED examined and not holding as stated).

Claim 1 — [CV]/[D]. The eigenvalue trinity (π,√5,φ) and the coupling g_c are the finite invariants of X(1): π from the area (3 vol=π), √5 from the discriminant (=τ(χ_5)), φ from the systole (4 log φ), g_c=1/(√5 π) as their forced combination.

Claim 2 — [D] (open reopener). ℚ(√5) is a richly-chosen, canonically-located marked point of X(1) — a forced home (the universality of X(1)) together with a chosen spot (the shortest geodesic) — not a point selected by a principle the framework supplies.

Claim 3 — [CV]. The sl(2,ℝ) commutator structure reaches the Selberg/Maass spectrum: the Casimir is the hyperbolic Laplacian, with spectrum $\lambda=¼+r^2$.

Claim 4 — [D]. The framework forces nothing beyond the single marked-point fact: the algebra is blind to the arithmetic content of the spectrum; the marked-point relation φ → systole is ceiling as well as floor; g_c is a product, not a bridge.

Claim 5 — [CV]. The framework's transfer operator is the Mayer operator; det(1−$\mathscr{L}$_s^2)=Z_Selberg(s) detects the zeros of the Selberg zeta — the Selberg/Maass spectrum — not the Riemann zeros.

Claim 6 — [D] account / [CV] constituents. The distinctive √5/φ content is confined to the length-spectrum channel and structurally barred from the ζ-governed scattering channel where the Riemann zeros live; the only arithmetic route carrying √5 into a zeta, ζ_ℚ(√5)=ζ· L(s,χ_5), changes the surface. The golden content is real but inert with respect to the ζ-zeros.

Claim 7 — [D]. The gap between this sound description and the Riemann Hypothesis is Weil positivity — which is to say RH itself; the "universal wall" is heuristic, not a no-go theorem.

Claim 8 — FIRED. The over-reaching layer of the framework — the claims that it forces the zeros, fixes external physical constants, or is independently confirmed — was examined and graded FIRED, and is excluded from this paper.

The whole characterization is held at [D] — a characterization consistent across independent analyses (strong corroboration, not a proof).

References

D. H. Mayer, "The thermodynamic formalism approach to Selberg's zeta function for PSL(2,ℤ)," Bulletin of the American Mathematical Society (New Series) 25 (1991), no. 1, 55–60.

J. B. Lewis and D. Zagier, "Period functions for Maass wave forms. I," Annals of Mathematics (2) 153 (2001), 191–258.

P. Sarnak, "Class numbers of indefinite binary quadratic forms," Journal of Number Theory 15 (1982), no. 2, 229–247.

C. Series, "The modular surface and continued fractions," Journal of the London Mathematical Society (2) 31 (1985), no. 1, 69–80.

H. Iwaniec, Spectral Methods of Automorphic Forms, 2nd ed., Graduate Studies in Mathematics 53, American Mathematical Society, Providence, RI, 2002.

D. A. Hejhal, The Selberg Trace Formula for PSL(2,ℝ), Vol. 1, Lecture Notes in Mathematics 548, Springer-Verlag, 1976; Vol. 2, Lecture Notes in Mathematics 1001, Springer-Verlag, 1983.

A. Terras, Harmonic Analysis on Symmetric Spaces and Applications, Vol. I, Springer-Verlag, 1985; Vol. II, Springer-Verlag, 1988.

Standard results invoked as established mathematics: the Gauss–Bonnet area of X(1); the factorization ζ_ℚ(√5)(s) = ζ(s)·L(s,χ_5) (the Dedekind zeta of the class-number-one field ℚ(√5)); and Weil's positivity criterion as an equivalent of the Riemann Hypothesis.