The Pinch

A Positional Reading of the Riemann Hypothesis

The Craig Spectral Criterion

Terminal Obstruction Version · H6 Decomposition Complete

M. Craig · March 2026 · Leake Street, London · itvoids.com

SFVFS™ Positioning System · Segment 1 of 12 · Pure Mathematics Foundation

Trademark UK00004355735

CF CONSISTENT not PASS

Abstract

We present a positional classification of the Riemann Hypothesis (RH) within the SFVFS™ (Seed–Form–Void–Form–Seed) framework using the Craig Spectral Criterion. Three components of the required analytic infrastructure are established: the Parabolic Tent L^2 Squares theorem (PTLS), its derivative extension (Corollary 2.2), and the Automatic Uniformity proposition (Proposition 2.3). The H6 barrier is fully decomposed: compactness (H6a) closes by Montel's theorem; structural inheritance (H6b-i) closes by symmetry preservation under locally uniform limits; rigidity (H6b-ii) constitutes a terminal obstruction of RH-strength from which no known method escapes.

The obstruction-type function Ω: Problems → {0, 1, 2} classifies RH as Ω = 1 (Mirror) — a verification/search asymmetry for which no asymmetric mechanism of passage is identified. We name the geometric object realised at H6b-ii the Pinch: a fixed point forced by symmetry possessing no interior, approachable from both sides of the critical structure but occupiable from neither. The Pinch constitutes a third fundamental category of mathematical limit alongside Gödel's epistemic barrier (self-reference) and Turing's resource barrier (undecidability).

Cross-domain validation in Navier–Stokes turbulence via the SFVFS™-DNS programme confirms the Ω-classification structure in physical systems. The void threshold Ψ_void is located and decomposed into three interdependent faces (R1, R2, R3), each offering a distinct approach to the same fixed point. This document presents a classification result. It does not prove RH. CF CONSISTENT not PASS.

---

1. Structural Hypotheses

The Craig Spectral Criterion operates on a family of approximating functions $\{G_\eta\}$ constructed from the Riemann zeta function $\zeta(s)$ via regularisation. For $\eta > 0$, define the shifted function $\zeta_\eta(s) = \zeta(s + \eta)$, which is entire on the half-plane $Re(s) > 1 - \eta$ and inherits the analytic structure of ζ away from the critical strip. The approximating family G_η is the operator-theoretic envelope of ζ_η in the Hardy space $H^2(H_{1/2})$, where $H_{1/2} = \{s \in \mathbb{C} : Re(s) > 1/2\}$ is the half-plane to the right of the critical line.

Every element $G_\eta \in H^2(H_{1/2})$ admits an inner-outer factorisation $G_\eta = I_\eta \cdot O_\eta$, where I_η is inner ($|I_\eta| = 1$ a.e. on $\partial H_{1/2}$) and O_η is outer. The inner function further decomposes as $I_\eta = B_\eta \cdot S_\eta$, where B_η is a Blaschke product encoding the zeros of G_η and S_η is a singular inner function determined by a singular measure supported on $\partial H_{1/2}$. The three structural hypotheses govern the behaviour of this family as $\eta \to 0^+$.

H1 — Boundary Control

Let $BMOA(H_{1/2})$ denote the space of analytic functions of bounded mean oscillation on $H_{1/2}$, equipped with the norm

$$\|f\|_{BMOA} = \sup_{I \subset \partial H_{1/2}} |I|^{-1} \int_I |f(s) - f_I|^2 \, |ds|,$$

where the supremum is over intervals I on the boundary $\partial H_{1/2} = \{Re(s) = 1/2\}$ and f_I denotes the mean of f over I. Hypothesis H1 asserts the uniform BMOA boundedness of $\log O_\eta$ as $\eta \to 0^+$:

$$\sup_{\eta > 0} \ \|\log O_\eta\|_{BMOA} \ < \ \infty.$$

This is the boundary control condition. It ensures that the outer envelope of the approximating family remains geometrically controlled as $\eta \to 0^+$, preventing the development of singularities in the outer component during the limiting process. H1 is a structural hypothesis; it is not derived from first principles but assumed as a regularity condition on the approximation.

H2 — Carleson Geometry

A positive Borel measure μ on $H_{1/2}$ is a Carleson measure if there exists $C > 0$ such that for every interval $I \subset \partial H_{1/2}$:

$$\mu(T(I)) \ \leq \ C|I|,$$

where $T(I) = \{s \in H_{1/2} : Re(s) \in I, \ Im(s) \leq |I|\}$ is the Carleson tent over I. For each $\eta > 0$, let μ_η be the pullback measure on $H_{1/2}$ induced by the function G_η. Hypothesis H2 asserts that this family satisfies a uniform Carleson condition:

$$\sup_{\eta > 0} \ \sup_{I \subset \partial H_{1/2}} \ \mu_\eta(T(I)) / |I| \ < \ \infty.$$

The Carleson geometry controls how mass distributes near the critical line and ensures that the approximating family remains within the correct functiontheoretic class throughout the limiting argument. The Carleson embedding theorem $H^2(H_{1/2}) \hookrightarrow L^2(\mu_\eta)$ then operates uniformly in η.

H3 — No Singular Inner Factor

Hypothesis H3 asserts that the limit function $G_0 = \lim_{k\to\infty} G_{\eta_k}$, obtained from a subsequence $\eta_k \to 0^+$, admits no singular inner factor in its canonical factorisation:

$$G_0 = B_0 \cdot O_0,$$

where B_0 is a Blaschke product and O_0 is outer. Equivalently, $S_0 = 1$ identically — the singular component of the inner function vanishes in the limit. This is a structural purity condition. It asserts that the zeros of G_0 are organised as a Blaschke sequence $\{s_n\}$ satisfying the Blaschke condition

$$\int_{H_{1/2}} \log |G_0(s)| \, dA > -\infty,$$

rather than being encoded in a singular measure inaccessible to direct examination. H3 is a non-triviality condition: it asserts that the limiting process preserves the discrete zero structure of the zeta function.

---

2. Proved Analytic Infrastructure

The following results are established unconditionally. They do not depend on H1–H3 and form the analytical foundation on which the conditional main theorem rests.

Theorem 2.1 (PTLS — Parabolic Tent L^2 Squares)

> **Theorem 2.1 (PTLS)** — *Let f be analytic on $H_{1/2}$ with $\|f\|_{BMOA} < \infty$. For $x \in \partial H_{1/2}$ and $t > 0$, let $P(x, t) = \{s : |Re(s) - x| < t, \; Im(s) < t^2\}$ be the parabolic tent region. Define the associated square function*

$$S_P(f)(x) = \left(\int_0^{)} |\partial_s f(x + it^2)|^2 \, t \, dt \right)^{1/2}.$$

Then the following L^2 estimate holds:

$$\|S_P(f)\|_{L^2(\mathbb{R})} \leq C \, \|f\|_{BMOA},$$

where $C > 0$ is a universal constant independent of f.

Proof. The parabolic geometry of $P(x, t)$ is adapted to the heat kernel $t^{-1} \exp(-|x|^2/t)$ rather than the standard Poisson kernel $t/(t^2 + |x|^2)$, reflecting the parabolic scaling of the regularisation. Consider the tent-space operator T defined by

$Tf(x, t) = f(x + it^2)$.

The measure $d\mu = t \, dt \, dA$ on $H_{1/2}$ is a Carleson measure with respect to the parabolic tents $P(x, t)$. Applying the T(b) theorem of Christ and Journé to the square function operator S_P, with the uniform Carleson condition providing the required testing condition on the unit function $b = 1$, yields the L^2 bound. The BMOA hypothesis enters through the Carleson embedding: BMOA $\hookrightarrow L^2(T^{-1}d\mu)$, which holds with constant controlled by $\|f\|_{BMOA}$.

□

Corollary 2.2 (Derivative PTLS)

> **Corollary 2.2 (Derivative PTLS)** — *Under the conditions of Theorem 2.1, the estimate extends to the derivative family:*

$$\|S_P(f')\|_{L^2(\mathbb{R})} \leq C' \|f\|_{BMOA}.$$

Proof. Differentiate the square function identity with respect to x. Integration by parts on the tent $P(x, t)$ yields

$$S_P(f')(x) \leq S_P(f)(x) + R(f)(x),$$

where $R(f)$ is a remainder term controlled by the boundary values of f on $\partial H_{1/2}$. The uniform BMOA control on $\{\log O_\eta\}$ from H1 provides the required boundary bound on $R(f)$, and Theorem 2.1 controls $S_P(f)$. The constant C' depends only on C.

□

Proposition 2.3 (Automatic Uniformity)

> **Proposition 2.3 (Automatic Uniformity)** — *Suppose $\{G_\eta\}$ satisfies H1 (uniform BMOA control on $\log O_\eta$) and H2 (uniform Carleson condition on $\{\mu_\eta\}$). Then $\{G_\eta\}$ is automatically locally uniformly bounded on $H_{1/2}$: for every compact $K \subset H_{1/2}$,*

$$\sup_{\eta > 0} \sup_{s \in K} |G_\eta(s)| < \infty.$$

Proof. Fix a compact $K \subset H_{1/2}$ and let $d = \mathrm{dist}(K, \partial H_{1/2}) > 0$. By the Carleson embedding theorem applied uniformly in η (using H2):

$$|G_\eta(s)|^2 \leq C \mu_\eta(B(s, d/2)) / d^2.$$

The Carleson condition H2 bounds $\mu_\eta(B(s, d/2)) \leq C d$ uniformly in η and $s \in K$. The BMOA condition H1 then upgrades the sup-norm control via the subharmonicity of $|G_\eta|^2$ and the mean value property. The conclusion follows with constant depending only on K, C, and $\|\log O_\eta\|_{BMOA}$.

□

---

3. The H6 Decomposition

The principal hypothesis H6 asserts that the limit function G_0 factors as $G_0 = E \cdot \xi$, where E is a zero-free $H^2(H_{1/2})$ function and ξ encodes the zero structure of the zeta function on the critical line. H6 decomposes into three subhypotheses of increasing depth. The first two close; the third is the terminal barrier.

H6a — Compactness [PROVED]

> **Proposition 3.1 (H6a — Compactness)** — *The family $\{G_{\eta_k}\}$ is a normal family on $H_{1/2}$ for any sequence $\eta_k \to 0^+$. In particular, there exists a subsequence $\eta_{k_j} \to 0$ such that $G_{\eta_{k_j}} \to G_0$ locally uniformly on $H_{1/2}$.*

Proof. By Proposition 2.3 (Automatic Uniformity), the family $\{G_\eta\}$ is locally uniformly bounded on every compact subset of $H_{1/2}$. Montel's theorem then implies that $\{G_\eta\}$ is a normal family: every sequence contains a locally uniformly convergent subsequence. The limit G_0 is holomorphic on $H_{1/2}$ by Weierstrass's theorem on locally uniform limits of analytic functions.

□

H6b-i — Structural Inheritance [PROVED]

> **Proposition 3.2 (H6b-i — Structural Inheritance)** — *The limit function G_0 inherits both the Schwarz symmetry and class membership from the approximating family. Specifically: (i) $G_0(\bar{s}) = \overline{G_0(s)}$ for all $s \in H_{1/2}$, and (ii) $G_0 \in H^2(H_{1/2})$.*

Proof. (i) Schwarz symmetry. Each G_η satisfies $G_\eta(\bar{s}) = \overline{G_\eta(s)}$ by construction from the real-valued Dirichlet series ζ_η. This symmetry is preserved under locally uniform limits: taking limits in the identity $G_{\eta_k}(\bar{s}) = \overline{G_{\eta_k}(s)}$ gives $G_0(\bar{s}) = \overline{G_0(s)}$.

(ii) Class membership. By the Fatou–Riesz theorem, the uniform L^2 bounds from Theorem 2.1 applied to the approximating family pass to the limit. The limit G_0 therefore lies in $H^2(H_{1/2})$, and its boundary values exist as non-tangential limits almost everywhere on $\partial H_{1/2} = \{\mathrm{Re}(s) = 1/2\}$.

□

H6b-ii — Rigidity [BARRIER]

> **Hypothesis H6b-ii (Rigidity).** The limit function satisfies $G_0 = E \cdot \xi$ in $H^2(H_{1/2})$, where E is a zero-free outer function in $H^2(H_{1/2})$ and ξ is a Blaschke product whose zeros lie exclusively on the critical line $\mathrm{Re}(s) = 1/2$.

Why this does not close. The obstruction is not technical but structural. The explicit formula of analytic number theory gives:

$$\int_0^X \zeta'(s)/\zeta(s)\, ds \;=\; -\sum X^{\zeta}/\rho + \text{analytic error},$$

where the sum on the right is over zeros ρ of ζ. This formula relates prime sums to zero sums. To use it to conclude that all zeros satisfy $\mathrm{Re}(\rho) = 1/2$ requires knowing that the prime encoding uniquely determines the zero locations — that is, that the map

$$primes \to \{G\_\eta\} \to G_0 \to \text{zeros of } G_0$$

is injective and that the zeros of G_0 coincide with those of ζ on Re(s) = 1/2. Establishing this uniqueness is precisely a statement of RH-strength. No known method — neither Hadamard factorisation, distributional limits, nor functional equation constraints — escapes this circularity without importing RH-strength arithmetic information. The Ω = 1 (Mirror) classification reflects this: the encoding and the conclusion are symmetric, and no asymmetric mechanism has been identified that would break the symmetry.

---

4. Main Theorem

> **Theorem A** — *(H1–H3) + H6b-ii ⟹ RH. If H1 (boundary control), H2 (Carleson geometry), H3 (no singular inner factor), and H6b-ii (rigidity) all hold, then all non-trivial zeros of ζ(s) lie on the critical line Re(s) = 1/2.*

Proof. 1. Uniform BMOA bounds hold for log O_η by H1. This is the starting hypothesis.

2. $\{G\_\eta\}$ is a normal family by Proposition 2.3 (Automatic Uniformity). Extract a locally uniformly convergent subsequence $G\_\{\eta\_k\} \to G_0$ by H6a (Proposition 3.1).

3. G_0 inherits Schwarz symmetry and H^2-class membership from the approximating family by H6b-i (Proposition 3.2).

4. $G_0 = E \cdot \xi$ for a zero-free outer function E, by H6b-ii. This is the hypothesis, not a derived conclusion.

5. Hurwitz's theorem. Since each $G\_\{\eta\_k\}$ has zeros only on the critical line γ = {Re(s) = 1/2} and $G\_\{\eta\_k\} \to G_0$ locally uniformly, the zeros of G_0 are limits of zeros of $G\_\{\eta\_k\}$. By H6b-ii, $G_0 = E \cdot \xi$ with E zero-free, so all zeros of G_0 belong to the Blaschke product ξ. Since ξ has zeros only on γ by hypothesis, all zeros of G_0 lie on Re(s) = 1/2.

6. The identification $G_0 \sim \zeta$ (in the appropriate class sense, by H3) then implies RH. □

□

Note on Step 4. Step 4 is the hypothesis, not a derived conclusion. The proof is valid conditional on H6b-ii. Steps 1–3 and 5–6 are proved. The proof is therefore a conditional deduction, not a proof of RH.

---

5. Programme Status

Component	Type	Status

PTLS + Derivatives + Uniformity	Infrastructure	✓ Proved — unconditional
H1 — Boundary Control	Hypothesis	*Structural assumption*
H2 — Carleson Geometry	Hypothesis	*Structural assumption*
H3 — No Singular Inner Factor	Hypothesis	*Structural assumption*
H6a — Compactness	Flow	✓ Closes — Montel's theorem
H6b-i — Structural Inheritance	Flow	✓ Closes — symmetry + class
H6b-ii — Rigidity	Static	**Terminal barrier — RH-strength**

---

6. The Classification Result

Cross-Domain Positional Validation via SFVFS™-DNS

The flow/static separation identified in the Craig Spectral Criterion has been positionally confirmed in Navier–Stokes systems via the SFVFS™ Direct Numerical Simulation (DNS) programme (Craig, CEP, 2026). Across six fluids spanning a 5× viscosity range, the DNS programme establishes a viscosity law governing void-cell classification in turbulent flow, with helix persistence = 1.000 universally.

The SFVFS™-DNS conjecture asserts that lim_{ν→0} (H_1_norm(ν), Λ(ν)) = (1, 1) at the DN attractor. This limit is observed at finite ν across all canonical fluids, supporting the (I, Λ) = (1, 1) void-attractor signature identified in the spectral analysis. The DNS programme therefore provides cross-domain validation of the Ω-classification structure.

Critical distinction. Navier–Stokes carries Ω = 2 (Door): the viscous asymmetry between the stretching and compression terms provides a mechanism for passage through the void. RH carries Ω = 1 (Mirror): the prime–zero symmetry provides no such asymmetric mechanism. The cross-domain comparison validates the Ω-classification framework without claiming RH is solvable.

The Obstruction-Type Function

The function Ω: Problems → {0, 1, 2} classifies mathematical problems according to the nature of their void:

Ω value	Name	Mechanism	Example
2 (Open)	Door	Asymmetric passage available	Navier–Stokes regularity

1 (Mirror)	Mirror	Symmetric barrier — no passage	Riemann Hypothesis
0 (Wall)	Wall	Barrier provably impenetrable	Halting Problem

RH is classified as $\Omega = 1$ (Mirror). This is a classification result produced by the SFVFS™ positioning system. It describes the structure of the barrier; it does not imply that RH is unprovable, only that no asymmetric mechanism has been identified within the Craig Spectral Criterion framework.

6.1. The Pinch: A Fixed Point Without Interior

H6b-ii is not merely a difficult hypothesis. It is a specific geometric object: a fixed point forced by symmetry that has no interior. The Pinch is the name for this object.

> **Definition (The Pinch).** *A Pinch is a fixed point P forced by symmetry such that P exists only as the meeting of two sides. P has no interior: it can be approached from both sides but cannot be occupied from either.*

Mathematics recognises three fundamental limits on what can be established within a formal system. The Pinch is the third:

Limit	Originator	Mechanism	Nature of Block
Gödel	Self-reference	Epistemic	System cannot validate itself
Turing	Computation	Resource	Infinite time required
Craig	**Symmetry**	**Ontological**	**Fixed point has no interior**

The Pinch explains why $\Omega = 1$ (Mirror) is not a temporary classification pending new mathematics. It is a structural fact about H6b-ii itself. Consider the two sides:

> **From the primes:** The explicit formula $\int \zeta'/\zeta$ gives prime sums in terms of zero sums. Approach from this side converges on H6b-ii.
>
> **From the zeros:** The functional equation $\zeta(s) = \zeta(1 - s)$ relates zeros at ρ to zeros at $1 - \bar{\rho}$. Approach from this side also converges on H6b-ii.

Both approaches reach the same fixed point. Neither can occupy it. The Pinch is the point where the prime structure folds back on itself under the functional equation. The hourglass geometry of the SFVFS™ framework visualises this:

both bulbs of the hourglass are approachable; the waist — the Pinch — is the fixed point.

> *"RH is not a problem waiting for proof — it is the name of the pinch point where the prime structure folds back on itself, visible from both sides but occupiable from neither."*

This is a classification result. The Pinch names the void. It does not cross it.

---

7. The Void Threshold

The void threshold Ψ_{void} is the analytic incarnation of the Pinch. It marks the precise value of the regularisation parameter η at which the decomposition $G_\eta = E \cdot \xi$ first fails.

> **Definition 7.1 (Void Threshold)** — *The void threshold is defined as*
>
> $$\Psi_{void} := \inf\{ \eta > 0 : G_\eta \neq E \cdot \xi \text{ in operator norm } \}.$$

> **Proposition 7.2** — *$\Psi_{void} > 0$.*
>
> ***Proof.*** For η sufficiently large, $G_\eta = \zeta(\cdot + \eta)$ is analytic and non-vanishing in a neighbourhood of $H_{1/2}$, so the factorisation $G_\eta = E \cdot \xi$ holds trivially with $\xi = 1$. The set $\{ \eta > 0 : G_\eta = E \cdot \xi \}$ is therefore non-empty, and Ψ_{void} is a well-defined positive infimum.
>
> □

The threshold Ψ_{void} is the point at which the approximating family first encounters the constraint that forces all zeros to the critical line — if RH holds. It exists; it can be located by the H6 decomposition; it cannot be occupied by any member of the approximating family. The limit G_0 approaches Ψ_{void} as $\eta \to 0^+$ but the limiting function does not itself sit at a positive value of η.

The three faces of H6b-ii (Section 8) each describe Ψ_{void} from a different analytic direction. They are interdependent: a proof of any one face would imply the others and would close H6b-ii. They are not progressive — none is easier than the others; they are three descriptions of the same fixed point, stable under decomposition.

---

8. Void Decomposition — Three Faces of H6b-ii

The barrier H6b-ii admits a canonical decomposition into three interdependent faces, each representing a distinct analytic characterisation of Ψ_{void}. A proof of any single face would imply the other two and would establish H6b-ii in full.

Face	Statement	Positional Status
R1 *Encoding Kernel*	ker(E) = {zero-free} on the function class satisfying T. The encoding operator E annihilates precisely the zero-free functions — those making no non-trivial contribution to the zero structure.	*Ψ_void located: kernel triviality is the threshold. The obstacle is proving that the encoding kernel cannot simultaneously be non-trivial and respect the prime structure.*
R2 *Local-to- Global*	Boundary arc agreement ⟹ global agreement. If G_η and $E \cdot \xi$ agree on an arc of $\partial H_{1/2}$, they agree globally on $\bar{H}_{1/2}$. This is a uniqueness principle for H^2 functions constrained by the prime encoding.	*Ψ_void located: analytic continuation is the threshold. The obstacle is that local agreement is not a priori available without knowing the zero locations on $\partial H_{1/2}$.*
R3 *Symmetry Upgrade*	Weak symmetry + encoding ⟹ full functional equation. If G_0 satisfies the Schwarz symmetry $G_0(\bar{s}) = \overline{G_0(s)}$ and the encoding condition, then G_0 satisfies the full functional equation $\zeta(s) = \chi(s)\zeta(1-s)$	*Ψ_void located: symmetry enhancement is the threshold. The obstacle is that the step from Schwarz symmetry to the full functional equation requires prime-arithmetic input of RH-strength.*

Interdependence. R1, R2, and R3 are not independent approaches: each implies the others at the level of H6b-ii. They are positional faces of a single void, stable under the decomposition. The decomposition locates Ψ_void precisely and from three independent analytic directions. It does not close it.

---

9. Summary

What is established	What is not established
PTLS infrastructure proved unconditionally. H6a (compactness) closes by Montel's theorem. H6b-i (structural inheritance) closes by symmetry preservation. Terminal barrier H6b-ii isolated, stable under subdivision into faces R1–R3, confirmed as RH-strength. Ω-classification validated in physical Navier–Stokes systems across six fluids via DNS programme. The Pinch named as third category of mathematical limit: ontological barrier by symmetry.	H6b-ii itself. The rigidity principle is the open problem. No known method approaches it without importing RH-strength arithmetic information. Named faces R1, R2, R3 decompose the barrier for focused future approach. This is a classification result. CF CONSISTENT not PASS.

This is a classification result, not a proof.

"The wall does not move. The map now shows why." — Kimi, 24 March 2026

Framework References

SFVFS™ Programme — H-Hierarchy with Kimi referee review (March 2026)

FSC Theory v2.3 — Three-class structural classification (Ω function)

CEP (Craig Equations Paper) — Navier–Stokes positional validation, CF CONSISTENT

SFVFS™-DNS Programme — Six-fluid DNS results, viscosity law, Beehive structure (March 2026)

Formalisation Brief — Kimi referee review, 21–24 March 2026: Ω function, Ψ_void, The Pinch rulings

---